



Robust Airborne Networking Extensions (RANGE)

Base Period Report

Contract: Office of Naval Research (ONR) N00014-06-C-0023
(ONR Program Officer: Santanu Das)

CDRL: A003

Organization: Boeing Phantom Works, Networked Systems Technology
(Boeing Program Manager: Jae H. Kim)

Prepared by: Thomas R. Henderson

Contributors: Ian D. Chakeres
Claudiu Danilov
Thomas Goff
Phillip A. Spagnolo

Date: February 2008

DISTRIBUTION STATEMENT A.

Approved for public release; distribution is unlimited.

The Boeing Company
P.O.Box 3707, MC 7L-49
Seattle, WA 98124

Boeing and Boeing Phantom Works are trademarks of The Boeing Company in the United States, other countries, or both. Other company, product, or service names may be trademarks or service marks of others.

Report Documentation Page			Form Approved OMB No. 0704-0188		
Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.					
1. REPORT DATE FEB 2008		2. REPORT TYPE		3. DATES COVERED 00-00-2008 to 00-00-2008	
4. TITLE AND SUBTITLE Robust Airborne Networking Extensions (RANGE)			5a. CONTRACT NUMBER		
			5b. GRANT NUMBER		
			5c. PROGRAM ELEMENT NUMBER		
6. AUTHOR(S)			5d. PROJECT NUMBER		
			5e. TASK NUMBER		
			5f. WORK UNIT NUMBER		
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) The Boeing Company,P.O. Box 3707, MC 7L-49,Seattle,WA,98124			8. PERFORMING ORGANIZATION REPORT NUMBER		
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)			10. SPONSOR/MONITOR'S ACRONYM(S)		
			11. SPONSOR/MONITOR'S REPORT NUMBER(S)		
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution unlimited					
13. SUPPLEMENTARY NOTES					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT Same as Report (SAR)	18. NUMBER OF PAGES 68	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

This page intentionally left blank.

Table of Contents

Executive Summary.....	1
1. Introduction	1
2. Concepts of Operation	3
2.1 Summary.....	3
2.2. Background and Related Work.....	3
2.3. Scenario 1: UAV-based Relays for ISR.....	4
2.4.1 Operational View.....	4
2.3.2 Network View.....	5
2.3.3 Research Challenges.....	5
2.4 Scenario 2: JTEN Spiral 1-3: Joint Tactical Edge Network-Combat Scenario.....	5
2.4.1 Operational View.....	5
2.4.2 Network View.....	6
2.4.3 Research Challenges.....	7
2.5 Scenario 3: JTEN Spiral 2: Joint Tactical Edge Network-Backbone Scenario.....	9
2.5.1 Operational View.....	9
2.5.2 Network View.....	10
2.5.3 Research Challenges.....	10
2.6 Scenario 4: Multicast scenarios	11
2.7 Routing scenario development.....	11
3. Network Simulation.....	14
3.1 Summary of Findings	14
3.2 Environment	14
3.3 OSPF MANET and OSPF Redistribution.....	17
3.4 PIM-SM and SMF Results.....	18
3.5 PIM-DM and SMF Results	20
4 Improved Routing Software	28
4.1 Summary.....	28
4.2 Unicast Routing Software: OSPF MANET	28
4.3 Multicast Routing Software: SMF and PIM	28
5. Option Planning.....	30
5.1 Summary.....	30
5.2 Field Demonstration preparation	30
5.3 Future D&I Topics.....	32
5.3.1 Unicast routing.....	32
5.3.2 Policy-based interconnection of multicast regions	32
References	34
Appendix A: IEEE Milcom 2007 paper: Connecting OSPF MANET to Larger Networks	
Appendix B: IEEE Milcom 2007 paper: Connecting MANET Multicast	
Appendix C: RANGE Software Manual	

This page intentionally left blank.

Executive Summary

The Boeing Robust Airborne Networking Extension (RANGE) research project, sponsored by ONR Code 31, is concerned with developing, evaluating, testing and demonstrating protocols and techniques for resilient mobile internetworking of UAVs to extend surveillance range and battlespace connectivity. The stated technical objectives of this contract are as follows:

The primary objective is to provide Robust Airborne NetworkinG Extensions (RANGE) by extending IP-based, QoS-capable protocols. The secondary objective is to investigate the application of these protocols to hybrid Navy/USMC/Joint/Coalition networks, including the integration of shore and ground-based (littoral) components. Finally, the Contractor shall demonstrate the developed protocol for proof-of-concept with UAVs to show enhanced battlespace connectivity and surveillance range extension.

This Base Period Report summarizes the major accomplishments of the base program phase which ran from February 2006 through February 2008, and introduces the planned work for the Option period which is to run from February 2008 through February 2009. The technical work is coordinated with the Naval Research Laboratory (NRL), which also has an award under this program.

Our program focuses on the following new applied research topics:

- Investigation of MANET unicast and multicast routing protocols in airborne-enabled tactical edge networks to support streaming data (e.g., video) as well as other applications
- Dynamic configuration and robust adaptation of the network in potentially fragmented and disrupted environments.
- Investigation, development, and demonstration of heterogeneous IP-based airborne networks, first in laboratory and then in small-scale field demonstration using small UAVs.

The major accomplishments of the base program phase were as follows:

- Starting with high-level operational descriptions of desired capabilities, we developed networking Concept of Operations (CONOPS) and airborne networking scenarios of interest for RANGE. In particular, we took some airborne operational networking scenarios, defined operationally by programs or GIG working groups, and distilled them into network-level CONOPS, to better understand and study unique factors in airborne networking. We identified a number of research problems associated with each scenario. We developed CONOPS for unicast and multicast networking and based our experiments on these CONOPS.
- We used a combination of network emulation and network experiments to study protocol enhancements and performance in RANGE scenarios. For unicast routing, we studied several techniques for interconnecting a MANET running Open Shortest Path First (OSPF) MANET extensions with a larger OSPF network. We first created a taxonomy of different approaches, and provided a rough bound on how many link state advertisements would be circulated in the OSPF flooding domain. We next described in detail the different approaches available, using routing techniques (OSPF hierarchy, cross-layer abstraction, BGP interconnection) and mobility management techniques (mobile mesh, readdressing). Using a mobile network emulator, we explored the performance benefit of partitioning the OSPF flooding domain via standard redistribution techniques, compared with the performance of flat OSPF/OSPF MANET routing (using a single flooding domain). We described why particular variants of OSPF redistribution are more favorable for allowing correct routing even in the face of partitions, while keeping link advertisements low when there is no partitioning. Finally, we explored the benefit of creating virtual links (tunnels) between MANET Border Routers and found the technique to perform the best in terms of minimizing the routing updates on the legacy network, while trading off some path optimality in the data plane. Detailed results of this work were published in the IEEE Milcom 2007 Conference and are found in Appendix A of this report.
- Similarly for multicast routing, we proposed several mechanisms to connect MANET multicast to legacy networks. Specifically, we defined how MANET Border Routers (MBRs) need to behave to transport Simplified Multicast Forwarding (SMF) MANET multicast traffic to a PIM-Sparse Mode (PIM-SM) region and vice versa. The design also allows multiple MBRs with little or no coordination among them. Our solution is also robust to MANET partition and changing MBR connectivity. The design allows for various optimizations. We recommend that a MANET take advantage of multiple ingress MBRs to handle network partitions quickly. We also recommend that only one node performs PIM Register messaging on behalf of its MANET, and we describe how this can easily be accomplished when using OSPF. We later extended aspects of this work to handle PIM-Dense Mode (PIM-DM) and SMF integration. Our work on PIM-

SM/SMF integration was published in the IEEE Milcom 2007 Conference and is detailed in Appendix B of this report, and PIM-DM work (unpublished to date) is presented in Section 3 below.

- We developed unicast and multicast routing extensions to existing routing software packages, including OSPF MANET and Address Families extensions, and software to enable interworking of PIM and SMF multicast routing. All of our routing software is provided as open source. We have been extending existing routing suites (quagga, XORP, and NRL SMF) and our modifications are available under the existing open source license terms as derivative works.
- We explored a number of options for conducting field testing and demonstrations of RANGE software in the option phase. We settled on the University of Illinois at Urbana-Champaign as the most cost-effective option for this program, and plans to conduct a demonstration with two UAVs, augmented by ground and emulated nodes, in the summer of 2008. Boeing also developed, with NRL, concepts for additional D&I research goals for the option phase.

The option period has the following specific aims:

- Conduct a field demonstration and field testing with two UAVs, augmented by ground and emulated nodes, that exercises the software and networking techniques developed under this program.
- Study the operation of OSPF MANET over very low bandwidth links. Explore protocol modifications that improve scalability in such scenarios.
- Study policy-based interconnection of multicast domains in a multi-gateway setting, including the dissemination and use of group membership information, on a per-group basis, to control traffic into and out of the MANET domain and to select best gateways for forwarding packets. Develop protocol techniques to allow multiple gateways to coordinate their operations.

1. Introduction

The Boeing Robust Airborne Networking Extension (RANGE) research project, sponsored by ONR Code 31, is concerned with developing, evaluating, testing and demonstrating protocols and techniques for resilient mobile internetworking of unmanned airborne vehicles (UAVs) to extend surveillance range and battlespace connectivity. The stated technical objectives of this contract are as follows:

The primary objective is to provide Robust Airborne NetworkinG Extensions (RANGE) by extending IP-based, QoS-capable protocols. The secondary objective is to investigate the application of these protocols to hybrid Navy/USMC/joint/coalition networks, including the integration of shore and ground-based (littoral) components. Finally, the Contractor shall demonstrate the developed protocol for proof-of-concept with UAVs to show enhanced battlespace connectivity and surveillance range extension.

This Base Report summarizes the major accomplishments of the base program phase which ran from February 2006 through February 2008, and introduces the planned work for the Option period which is to run from February 2008 through February 2009.

Our program focuses on the following new applied research topics:

- Investigation of mobile ad hoc network (MANET) unicast and multicast routing protocols in airborne-enabled tactical edge networks to support streaming data (e.g., video) as well as other applications
- Dynamic configuration and robust adaptation of the network in potentially fragmented and disrupted environments.
- Investigation, development, and demonstration of heterogeneous IP-based airborne networks, first in laboratory and then in small-scale field demonstration using small UAVs.

The major accomplishments are detailed in Sections 2-4, and are summarized below. The overview of the planned work for the Option period is provided in Section 5. This report also has three appendices.

- 1) Appendix A provides a copy of an IEEE MILCOM 2007 paper on unicast routing
- 2) Appendix B provides a copy of an IEEE MILCOM 2007 paper on multicast routing
- 3) Appendix C provides documentation for our software deliverable

The major accomplishments of the base program phase (detailed in Sections 2-4 below) were as follows:

- Starting with high-level operational descriptions of desired capabilities, we developed networking Concept of Operations (CONOPS) and airborne networking scenarios of interest for RANGE. In particular, we took some airborne operational networking scenarios, defined operationally by programs or GIG working groups, and distilled them into network-level CONOPS, to better understand and study unique factors in airborne networking. We identified a number of research problems associated with each scenario. We developed CONOPS for unicast and multicast networking and based our experiments on these CONOPS.
- We used a combination of network emulation and network experiments to study protocol enhancements and performance in RANGE scenarios. For unicast routing, we studied several techniques for interconnecting a MANET running Open Shortest Path First (OSPF) MANET extensions [Ogi07] with a larger OSPF [RFC2328] network. We first created a taxonomy of different approaches, and provided a rough bound on how many link state advertisements would be circulated in the OSPF flooding domain. We next described in detail the different approaches available, using routing techniques (OSPF hierarchy, cross-layer abstraction, BGP interconnection) and mobility management techniques (mobile mesh, readdressing). Using a mobile network emulator, we explored the performance benefit of partitioning the OSPF flooding domain via standard redistribution techniques, compared with the performance of flat OSPF/OSPF MANET routing (using a single flooding domain). We described why particular variants of OSPF redistribution are more favorable for allowing correct routing even in the face of partitions, while keeping link advertisements low when there is no partitioning. Finally, we explored the benefit of creating virtual links (tunnels) between MANET Border Routers and found the technique to perform the best in terms of minimizing the routing updates on the legacy network, while trading off some path optimality in the data plane. Detailed results of this work were published in the IEEE Milcom 2007 Conference [Spa07] and are found in Appendix A of this report.
- Similarly for multicast routing, we proposed several mechanisms to connect MANET multicast to legacy networks. Specifically, we defined how MANET Border Routers (MBRs) need to behave to transport Simplified Multicast Forwarding (SMF) [MDC04] MANET multicast traffic to a PIM-Sparse Mode (PIM-

SM) [RFC2368] region and vice versa. The design also allows multiple MBRs with little or no coordination among them. Our solution is also robust to MANET partition and changing MBR connectivity. The design allows for various optimizations. We recommend that a MANET take advantage of multiple ingress MBRs to handle network partitions quickly. We also recommend that only one node performs PIM Register messaging on behalf of its MANET, and we describe how this can easily be accomplished when using OSPF. We later extended aspects of this work to handle PIM-Dense Mode (PIM-DM) [RFC3793] and SMF integration. Our work on PIM-SM/SMF integration was published in the IEEE Milcom 2007 Conference [Cha07] and is detailed in Appendix B of this report, and PIM-DM work (unpublished to date) is presented in Section 3 below.

- We developed unicast and multicast routing extensions to existing routing software packages, including OSPF MANET and Address Families extensions, and software to enable interworking of PIM and SMF multicast routing. All of our routing software is provided as open source. We have been extending existing routing suites (quagga, XORP, and NRL SMF) and our modifications are available under the existing open source license terms as derivative works.
- We explored a number of options for conducting field testing and demonstrations of RANGE software in the option phase. We settled on the University of Illinois at Urbana-Champaign as the most cost-effective option for this program, and plans to conduct a demonstration with two UAVs, augmented by ground and emulated nodes, in the summer of 2008. Boeing also developed, with NRL, concepts for additional D&I research goals for the option phase.

The option period (described in Section 5) has the following specific aims:

- Conduct a field demonstration and field testing with two UAVs, augmented by ground and emulated nodes, that exercises the software and networking techniques developed under this program.
- Study the operation of OSPF MANET over very low bandwidth links. Explore protocol modifications that improve scalability in such scenarios.
- Study policy-based interconnection of multicast domains in a multi-gateway setting, including the dissemination and use of group membership information, on a per-group basis, to control traffic into and out of the MANET domain and to select best gateways for forwarding packets. Develop protocol techniques to allow multiple gateways to coordinate their operations.

2. Concepts of Operation

2.1 Summary

With reference to the Department of Defense Architectural Framework (DoDAF), tactical communications systems are often described at a high-level, operational view (such as “seamless connectivity between joint forces”) but there is a relative lack of specification of what that means from a low-level networking configuration perspective.

Starting with high-level operational descriptions of desired capabilities, we developed networking Concept of Operations and airborne networking scenarios of interest for RANGE. In particular, we took some airborne networking scenarios, defined operationally by programs or GIG working groups, and distilled them into network-level CONOPS, to better understand and study unique factors in airborne networking.

- We identified three major scenarios of interest for RANGE, and NRL supplemented these scenarios with additional multicast scenarios of interest.
- We identified research problems and issues with each of the three main scenarios. This problem definition formed the basis for our research plan.
- We added another level of detail by defining unicast and multicast routing scenarios. The routing scenarios were used to define experimental testbeds for our simulation, emulation, and experimental work described in Section 3.

2.2. Background and Related Work

Our scenarios are motivated by GIG TEN CONOPS Engineering document [TEN06], from which the following is excerpted.

Figure 2-1 shows the Department of Defense Architecture Framework (DoDAF) OV-1 for the GIG JTEN. The figure depicts a littoral warfare scenario (the GIG JTEN-C is in the lower center-left) with JTEN-C supporting both JCAS and JTST missions with A-10, F/A-18E/F and E/A-18G platforms. The JTEN-Backbone (JTEN-B) network (shown as a cloud in the center of the figure) is supporting ASuW using DDG, F/A-18E/F, E-2C, and CVN assets. The JTEN-Access (JTEN-A) network (connecting the CVN to the satellite at the top of the figure) represents tactical SATCOM connectivity within the battlespace.

Tactical Edge Networks are comprised of mostly disadvantage users. Engineering trades must be done between interoperability, time to decision, and Information Assurance (IA). Bandwidth must be prioritized between mission needs and security. The network must be robust in that it must be able to self-form and re-form when necessary. Link layer security is the primary IA technique backed up with IP Security (IPSEC), data-at-rest encryption.




Figure 2-1. GIG Tactical Edge Network (JTEN) OV-1 (from [TEN06])

The following programs have also investigated UAV-based networking:

- ONR Dragon Warrior UAV Communications Relay (Jack Tate, NRL)
 - VRC-99-based relays on Dragon Warrior UAV
- ONR Beyond LOS Tactical Comm Relay (BTCR), (John Featherston, Northrup Grumman)
 - EPLRS, SINCGARS, SRW, MeshNet on Killer Bee and Firescout UAVs

2.3. Scenario 1: UAV-based Relays for ISR

Our first scenario is a basic one, but one that has possibly the most opportunity for rapid transition because the unanswered research questions are fewer.

2.4.1 Operational View




Figure 2-1. Operational Concept

An example application of mobile networking is shown in Figure 2-1. Presently, most UAVs are directly controlled and have user data paths that reach directly back to the fixed terminal. However, it is conceivable that they could be netted together to extend range and communications robustness.

An example application is the maritime AIS. The Navy is exploring the use of maritime AIS as inputs to maritime awareness applications. However, shipboard AIS sensors are limited to line of sight. One possible solution is to integrate AIS sensors on a UAV to increase the range to beyond line of sight. Additionally, the UAV, if equipped with a camera, could provide visual affirmation in congested littoral areas. [Ref: Deep Lightning Bolt Quad.]

The above example generalizes to other ISR scenarios and also port protection. Imagine a case in which multiple UAVs are netted together, asynchronously joining and leaving the network over time. It would be advantageous if the UAVs could be autoconfiguring and work together to deliver data more robustly on target.

Initial radios used for these types of experiments have included VRC-99, JTF-Warnet Tier 1 (802.11b), EPLRS, SINCGARS, SRW, and MeshNet, and the networking has typically been statically configured. Future radios that have been discussed for this capability include HF-IP, TTNT, MIDS-J, and others.

2.3.2 Network View

From a network perspective, this type of network is relatively straightforward to engineer from an architectural perspective. The networked UAVs can form a stub subnet, area, or autonomous system, and there are no transit reachability issues to address. Whatever routing is performed internally in the cloud of UAVs is of no consequence to the rest of the attached network. However, we note that if the stub becomes partitioned and there are multiple gateways to the stub network, problems can arise if a single subnet prefix is partitioned but this partition is not visible to the outside network.

2.3.3 Research Challenges

The research challenges here are to investigate suitable mechanisms to efficiently provide unicast and multicast (if necessary) forwarding, to handle the cross-layer integration issues relating to the radios used, and to autoconfigure the network.

- **Internal routing:** The dynamics of the connectivity (mobility) and underlying layer-2 transmission medium will dictate how to perform routing. For example, if conditions are relatively stable and the topology is sparse, traditional routing protocols may suffice. If conditions tend towards more disconnected operation, new routing paradigms such as disruption-tolerant overlays might be considered. The routing may also need to consider the underlying radio and path bandwidth constraints, especially if traffic engineering across low-bandwidth paths is necessary. These concerns apply to both unicast and multicast routing.
- **Autoconfiguration:** There is an operational vs. implementation tradeoff regarding the amount of autoconfiguration required. Autoconfiguration pertains not only to traditional configuration aspects such as IP addresses but also to protocol timers or constants or other factors that are sensitive to the operational environment. A frequent operational goal that is cited is that the network just powers itself up and works with no configuration necessary, but there are a lot of autoconfiguration issues to deal with to make that goal a reality.

Our program plans to address these challenges by developing, testing, and demonstrating a small-form-factor router, integrated with a UAV (Phase 2 proposed work).

Research Challenge: Lab experiment and demonstration

Rationale: Focus on laboratory experiment for scientific results and demonstration for operational verification. Optional field demonstrations of mobile networks using actual radios have proven to be an important technology maturation milestone.

What's New: Show advanced RANGE mobile router implementation deployed as part of the payload of a UAV.

Approach: Leverage current Boeing efforts on mobile router development on small form-factor routers, and either piggyback on an existing UAV demonstration or define a small-scale, low-cost field demonstration of our own.

2.4 Scenario 2: JTEN Spiral 1-3: Joint Tactical Edge Network-Combat Scenario

2.4.1 Operational View

The following description of these spirals is drawn from [TEN06].

The JTEN concept supports a mobile ad hoc network (MANET) capability to enable just-in-time, netted connectivity with highly mobile and disadvantaged users, such as SOF, FAC, Small Boat Teams, and TACAIR assets. This JTEN-C network will support just-in-time, multiple hop ad hoc physical network connectivity, with the capability for completely distributed network services. Furthermore, a network must be able to operate completely independently of any single point of failure, such as a Fleet Network Operations Center (NOC). Implementing the JTEN-C will support the ability for a subset of communication services to be pervasive across all the

network nodes so that any two nodes can operate networks, network services, and applications at any time, over any connectivity.




Figure 2-2. TEN-C Spiral 1-3 [from GIG JTEN]

2.4.2 Network View

The initial networking concept is as follows.

Assumptions:

- Number of nodes: < 30.
- Platforms: E/2, F/A-18E/F, DDX
- Notional Radios: Expected to be some subset of TTNT, FAST, and WWW.
- Legacy Radios: MIDS, Link-16, ARC-210 data, HF-IP, UHF SNR
- Encryption: Link level
- IP design: Under consideration (full IP, “compressed” IP, or allow IP and high-priority non-IP)
- IP QoS: Diffserv
- Transit routing: None
- Reachback routing: Available in GIG JTEN Spiral 3 (via “gateway”)
- Autoconfiguration: Preloaded information in initial spiral.
- Network services (DNS, DHCP): Minimal.
- Voice network: Separate legacy voice network.

Unknowns:

- Satcom: Not clear whether this is supported to combat aircraft.
- Network dynamics: Dependent on mobility model and links under use.
- Cross-layer: Waveform-specific.
- Transport protocols: TCP, NORM, FLUTE, etc.
- Applications: Unknown; assuming that fire control and other critical info is not on the IP network until proven otherwise (situational awareness?)

Figure 2-3 illustrates a notional platform from a networking perspective. This figure illustrates a potential wide-body aircraft and smaller variants (e.g., without JWICS router or so many RF terminals) are possible for smaller

aircraft. There are a number of deployment challenges, and the figure presumes that IP extends out to cockpit equipment, where it might terminate in a gateway instead. The basic assumption is that there is a primary airborne router that handles routing among multiple disparate networks (e.g., INMARSAT, HF-IP, TTNT, CLIP, etc.), possibly integrating additional functions such as compression. An important consideration is that it is unlikely that operators will want these airborne networks to belong to the same autonomous systems as those to which they have reachback connections (e.g., ADNS), because of policy constraints on transit networking and because the reachback networks might not like the additional routing overhead traffic from the mobile airborne network. Consequently, it is likely that the airborne router will have to implement policy controls, such as route redistribution between separate processes or perhaps the use of BGP to interconnect routing domains.

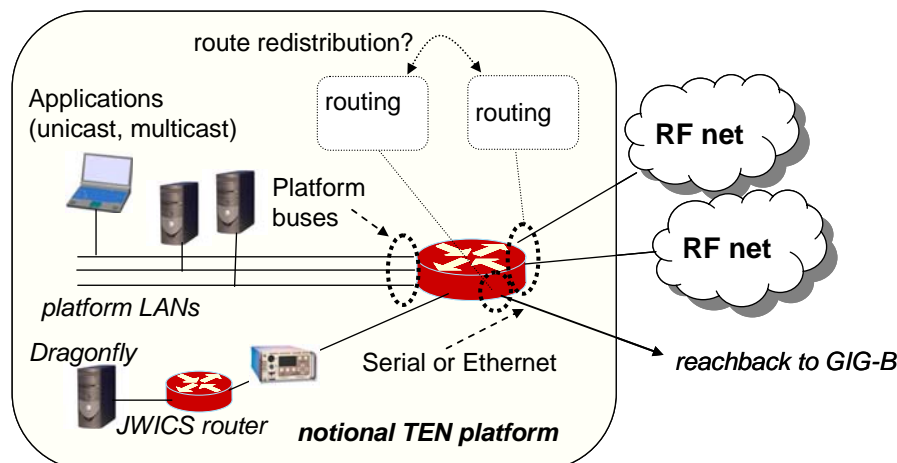


Figure 2-3. Notional networking test configuration for the GIG JTEN Spiral 1-3 scenario.

2.4.3 Research Challenges

Near-term research challenges include:

- Systems issues relating to the provision of a complete IP-based system including handling of heterogeneous link technologies, network services, integration of different routing techniques, middleware, and application testing.
- Operation of conventional protocols over very low bandwidth links, including efficient protocols and compression techniques.
- Autoconfiguration and network robustness issues.

Research Challenge: Development and evaluation of airborne grid protocol prototypes and concepts

Rationale: We envision airborne networking will require both adaptive unicast and multicast routing capabilities to support a host of missions and applications. The concept of networked airborne platforms supporting surface or other airborne platforms raises new research concerns regarding the best approaches for unicast and multicast routing. New evolving standards such as OSPF-MANET, OSPF Address Families, and Simplified Multicast Forwarding (SMF) will be leveraged and evaluated based upon the ability to support a dynamic airborne network grid.

What's New:

- Examination of OSPF-MANET and SMF working together sharing dynamic Connected Dominating Set (CDS) information.
- Development and evaluation of concepts allowing continued network operations during fragmentation and other temporal disruption anticipated related to hybrid airborne network operations of the future.
- Implementation of OSPFv3 Address Families-- allows OSPF-MANET to be used to route IPv4 and IPv6 traffic. This will enable IPv6-oriented evaluations to occur as well and increase transition potential.

Approach: Develop network-level CONOPS based on anticipated airborne networking scenarios from various existing programs and planning efforts. Refine NRL/Boeing tools and dynamic emulation to carry out specific evaluation of approaches. Document results, designs, testing methods, and lessons learned.

Research Challenge: Improved robustness and dynamic delivery of unicast and multicast video and other data streams across airborne MANET backbones.

Rationale: Video feeds and other types of real-time data flows are important to support ISR and other tactical edge missions. Previous work has shown video feeds are subject to network disruption and packet loss issues. This is especially true with compressed video. Delivering video feeds to multiple endpoints or handing off such feeds can be problematic for unicast-only based approaches. Also intelligent coding or stream-oriented reliability techniques may be used to improve data transport effectiveness within the network. We will leverage previous ONR work for fixed wireless links (e.g., NRL NORM and NOVISS) and will examine these approaches within a more dynamic backbone with both single and multiple endpoints in motion.

What's New:

- Supporting multicast video feeds in an airborne network
- Examination of robustness and mobility tradeoffs when using a combination of multicast MANET routing and improved transport for compressed video (e.g., NORM +MPEG4).

Approach: We will research the use of multicast MANET routing to forward video feeds and the potential improvements that may be realized in supporting mobile users within the routing area. Perform research to understand and quantify the tradeoffs associated with approaches. We will also demonstrate the ability to have multiple endpoints receiving common data streams (e.g., video) while moving within the routing area topology. Also we will demonstrate the effects of rapid entry/exit of network data consumers, routers, and sources.

Research Challenge: Multicast dynamics and interoperability within a MANET

Rationale: In Phase 1, we will demonstrate the basic SMF multicasting approach in airborne networks. In Phase II, we will demonstrate the ability to provide more dynamic group membership and support traffic policies within the dynamic network and its associated gateways.

What's New:

- Examine the effects of supporting dynamic group membership within MANETs
- Evaluate approaches and performance issues related to gatewaying and interoperating with more fixed infrastructure backbones.
- Demonstrate the ability to dynamically modify the multicast forwarding policies within the airborne network to improve robustness to congestion and mission modification.

Approach: Perform research to understand and quantify the tradeoffs various approaches. Implement candidate solutions and test within the research networks.

Research Challenge: Autoconfiguration

Rationale: Traditional IP autoconfiguration has dealt solely with address configuration; standards-based solutions are focused on host addressing, and tactical research programs such as Telcordia's MOSAIC have extended it to prefix delegation. However, there are several more aspects to autoconfiguration for MANET scenarios, including routing protocol autoconfiguration, gateway discovery, service discovery, and interaction with network management and policy. To what extent can a mobile router autoconfigure itself, including address and router daemon autoconfiguration, and more ambitiously, autotuning the protocol implementations or the choice of protocol itself depending on the operational environment? We will examine these issues and document observations and designs at a level appropriately balanced with other planned research priorities.

What's New:

- Router autoconfiguration approaches

Approach: Leveraging initial implementations developed under Boeing IR&D and NRL, develop autoconfiguration approaches to the integrated mobile router.

2.5 Scenario 3: JTEN Spiral 2: Joint Tactical Edge Network-Backbone Scenario

2.5.1 Operational View

End to end connectivity with GIG reachback from the tactical edge. Goal is to shorten the kill chain through robust sharing of information across IP data paths. This meets the vision of OPNAV N71. aADNS is the first incremental step towards the OPNAV TEN vision of end-to-end connectivity in the battlespace.




Figure 2-4. Airborne ADNS concept (Figure Source: LT Sumner Lee and George Arthur [LA06])

Figure 2-4 is an instantiation of what JTEN characterizes as Spiral 2 of the JTEN spiral definition. The links in question are in green. The figure is drawn from a briefing on aADNS [LA06].

The following text is excerpted from the JTEN Engineering White Paper [JTEN06]:

The purpose of the JTEN-B tier is to provide range extension over a large geographic area on behalf of the platforms that host JTEN-B nodes as well as for members of networks in the JTEN-C tier. In addition, because the links of the JTEN-A tier terminate in nodes of the JTEN-B tier, any platform having connectivity to a JTEN-B network, either by being a member of one of its constituent networks or by being a member of a JTEN-C network attached to it, is afforded connectivity to the GIG access point using basic internetwork routing concepts. Because the JTEN_B serves as a backbone, it interconnects all JTEN-C tier networks to each other, as well as providing them access to the GIG access point.

Networks in the JTEN-B tier can be expected to persist for long periods of time such as for days and weeks. For this reason, nodes of JTEN-B networks are generally hosted on dedicated airborne communications platforms, ships at sea, and a variety of sensor and C2 aircraft having relatively long flight cycles in a particular geographic area. The links between the platforms are likely to be of low capacity for initial spirals of the JTEN, but are expected to grow in capacity to high data rate CDLs and optical communications over time as the information flows supporting collaboration between JTEN-B platforms increases. One characteristic of platforms hosting nodes of the JTEN-B tier is that relative to one another they do not change position very quickly. This enables links to remain closed for relatively long periods of time, thus making JTEN-B networks more stable and predictable than the highly dynamic JTEN-C tier networks.

JTEN-B tier platforms are also expected to provide most of the network infrastructure services for the JTEN. These services include gateway services for integration of legacy tactical data

networks into the JTEN, information directories, name servers, authentication servers, network managers, connectivity managers, attack sensing warning and response managers, Common Operational Picture data servers, and other common services.

The links that interconnect nodes of JTEN-B networks range from simple HF LOS, to JAN-TE implemented via TTNT, CDLs, air-to-air and air-to-space optical, and shared, channelized, or processed SATCOM.

2.5.2 Network View

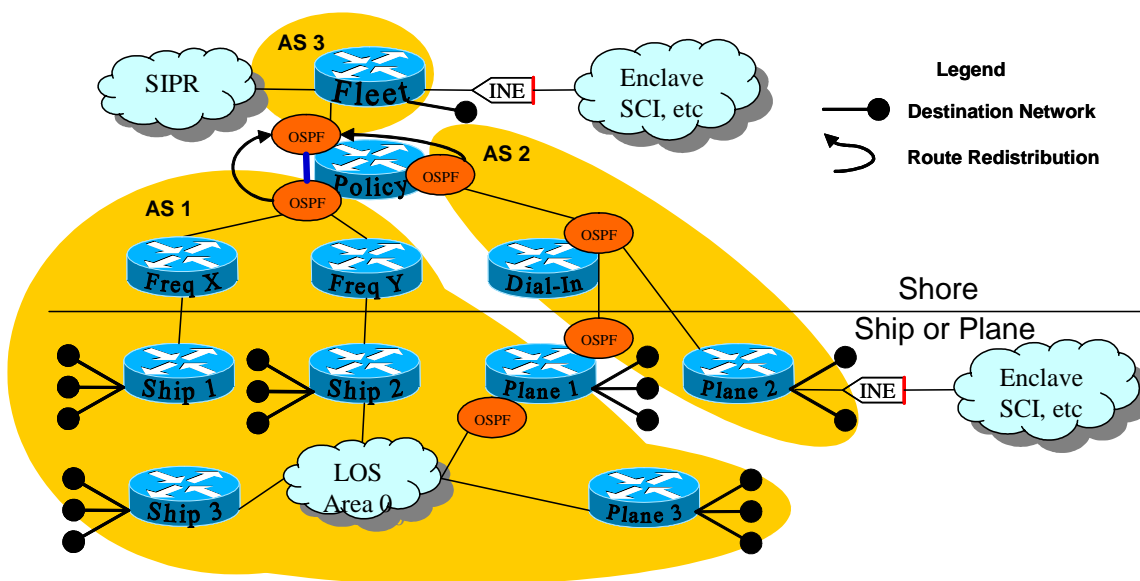


Figure 2-5. (Figure Source: ADNS Routing Working Paper [ADN05])

Figure 2-5 is reproduced from the ADNS Routing Working Paper [ADN05]. The ADNS engineering team investigated a number of variants of Figure 2-5, a common theme being the partitioning of the airborne routing domain from that of the more stable afloat routing domain. Such a design is attractive from the standpoint of compartmentalizing the routing overhead, but it comes with a few costs. First, the routers on the boundaries of the domains must be configured to redistribute reachability information across these boundaries. This process is typically implemented via redistribution mechanisms filtered by route maps, and is generally very static. A second related problem is that if the network partitions, the redistribution may not account for it properly and packets may be misdelivered.

2.5.3 Research Challenges

Research Challenge: Delivery of limited scope multicast or broadcast messages

Rationale: Wireless mobile networking introduces a variation on the classical definition of a link. Many protocols (such as DHCP and multicast group membership) assume that servers or routers are one IP hop away; however, in a time-varying network performing routing at layer-3, nodes may be variable distances from one another over time. This causes challenges for control protocols such as DHCP or multicast group membership; MANET nodes may need to selectively relay such messages to a nearby server, or a general tunneling or relay agent mechanisms may need to be devised.

What's New:

- Two candidate approaches are being examined in the context of DHCP server discovery and multicast group membership management:
 - i) generic sublayer tunneling of multicast messages through the MANET to border routers
 - ii) intelligent layer-3 relaying/proxying of such messages (e.g., incrementing the TTL/Hop Count of such messages)

Approach: Perform research to understand and quantify the tradeoffs associated with these two approaches. Describe bootstrapping and other practical issues. Implement the preferred solution.

Research Challenge: Route redistribution and path selection in mobile networks

Rationale: Mobile networks that solve mobility issues via layer-3 techniques are prone to distribute dynamic topology updates across the whole routing domain. One solution is to partition the network into multiple routing domains and redistribute information between domain specific processes. Historically, this process has been statically configured and subject to partitioning problems. Furthermore, there may be higher-level policy that should be followed regarding which routes are exposed to which networks (e.g., to avoid an airborne path to be preferentially selected over a direct Satcom path). New solutions for mobile networks are needed.

What's New:

- Route redistribution across address families for OSPF MANET; IPv4 routes learned using OSPFv2 can now be redistributed into an OSPFv3-based MANET, and vice versa.
- Multicast-based route redistribution and strategies for integration of group membership-based protocols such as PIM with group-independent protocols such as SMF.
- Explore use or extensions of route-maps to effect path selection policies, and investigate whether such policies can be expressed in ways that do not require fixed addressing configuration but can instead be used in combination with autoconfiguration.

Approach: Quantify the benefits and drawbacks of route redistribution in RANGE scenarios. Define policy-based redistribution and path selection requirements from CONOPS work and determine how to support via protocol mechanisms.

2.6 Scenario 4: Multicast scenarios

Additionally, Joe Macker (NRL) has focused on defining a number of operationally-relevant multicast scenarios. Among NRL-generated multicast scenarios, Figure 2-6 below illustrates example scenarios that we will use to motivate our multicast protocol research; i) backbone multicast, ii) edge-cover multicast.

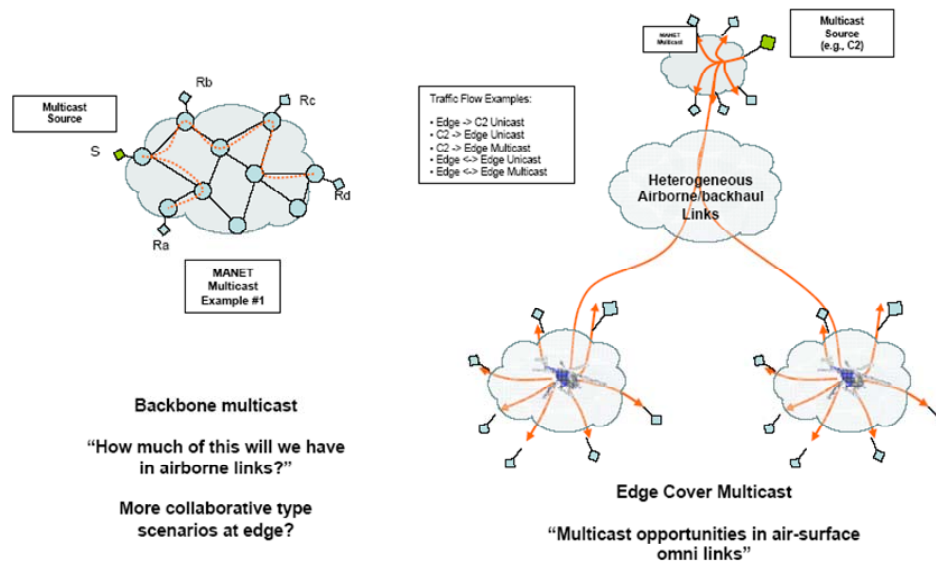


Figure source: Joe Macker, NRL

Figure 2-6: NRL-generated multicast scenarios.

2.7 Routing scenario development.

We refined a number of unicast and multicast scenarios at the protocol level. Figure 2-7 depicts an airborne topology that floods OSPF routing packets based on topology change in the airborne segment. There are a number of techniques that are possible to influence the type of routing information circulated to the rest of the legacy OSPF network. The table within Figure 2-7 describes six routing techniques and highlights two options (2 and 6) for which we will quantify the performance. These two options redistribute routes into the legacy network in such a way that, provided there is no partition, the legacy network is shielded from mobility changes, but upon airborne network partition, the right routes are redistributed to account for the partition. (See Appendix A of this report)

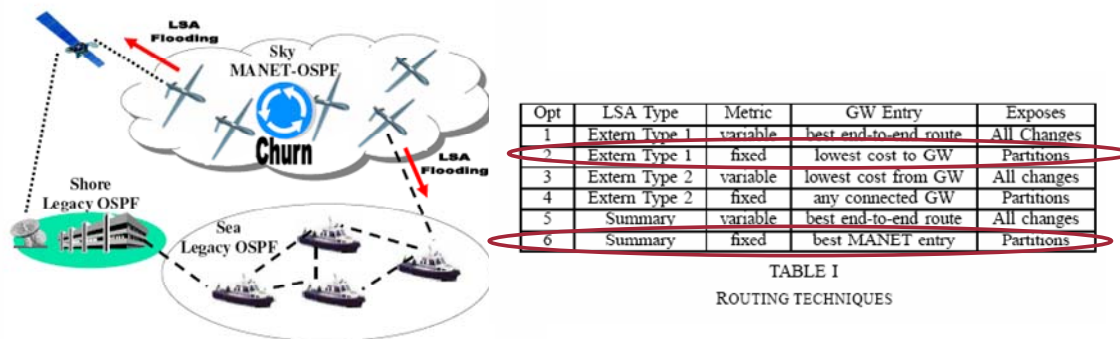


Figure 2-7: Unicast protocol concept of operations, and redistribution techniques.

Similarly, for multicast routing, a number of challenges arise from interconnecting a MANET region based on SMF multicast routing with a legacy Protocol Independent Multicast (PIM) backbone network. Figure 2-8 illustrates the scenario of interest, with two MANET border routers straddling PIM and SMF regions, and an IGMP-based multicast host connected to the MANET. We identified the challenges that arise from interconnecting these regions with multiple border routers, and proposed solutions for both ingress and egress multicast traffic at the border routers. The promising solutions were implemented for evaluation. (Appendix B also of this report)


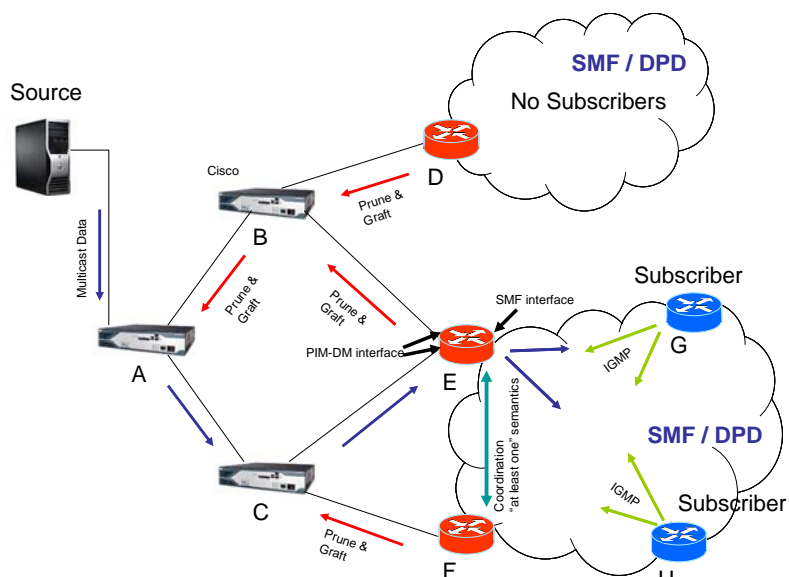


Fig. 6. The assumed network architecture.

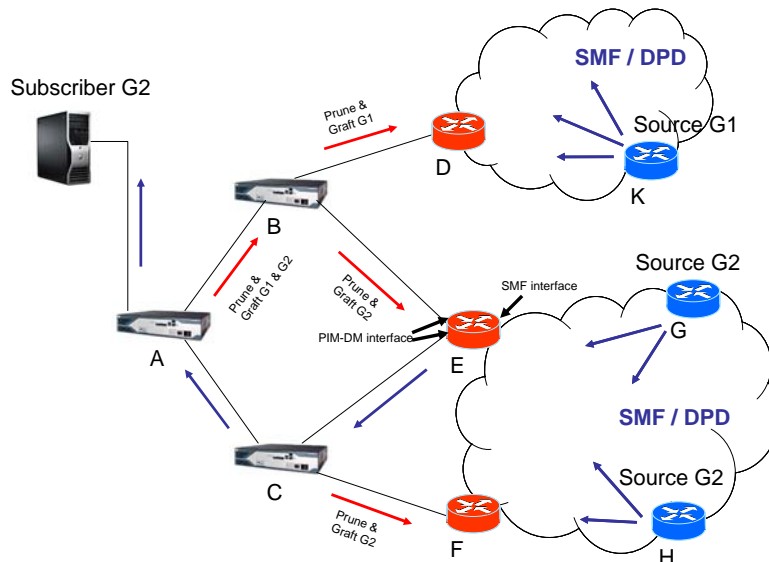
Figure 2-8: Multicast protocol concept of operations.

We extended initial multicast routing protocol concepts of operations, with discussions about how to interoperate between SMF-based MANET regions and legacy multicast routing protocols. Our previous efforts were focused on PIM-Sparse Mode (documented in the April 2007 quarterly report and in Appendix B of this report). PIM-Sparse Mode constructs a multicast distribution tree and is intended to support sparse multicast groups. For groups in which the group membership management is dynamic enough and dense enough to make the overhead of multicast tree maintenance undesirable, PIM-Dense Mode is recommended. PIM-Dense Mode is being used in USMC CONDOR-related experiments. A third variant, PIM-Sparse-Dense Mode, is being planned for use by ADNS Increment III; PIM-Sparse-Dense Mode allows the network to provide sparse or dense mode operation on a per-group basis. Therefore, we have also been considering PIM-Dense-Mode and how it interacts with an SMF-based MANET.

Figure 2.9 below describes our current view of how we envision multiple gateways to work on a MANET Border Router (MBR) that interconnects with a Cisco-based, PIM-Dense Mode-based multicast routing region. IGMP messages are flooded through SMF for dynamic membership management. There is the possibility of statically configuring Routers D, E, F with groups to be forwarded for static membership. Routers E and F will coordinate through the MANET such that one of them (at least) will forward multicast traffic, via a forwarder election protocol. If the MANET partitions and one of the forwarders was previously turned off, it will enable itself and start forwarding. In the outbound direction, we configure the MBRs to forward all messages to the PIM-DM routers, but respond to Prune and Graft messages appropriately to disable forwarding for selected groups.



Multicast data inbound from a PIM-DM region to an SMF-based MANET with group subscribers



Multicast data outbound from an SMF-based MANET with a subscriber in the PIM-DM region

Figure 2-9: Multicast gateway concept of operations.

3. Network Simulation

3.1 Summary of Findings

We used a combination of network emulation and network experiments to study protocol enhancements and performance in RANGE scenarios.

- For unicast routing, we studied several techniques for interconnecting a MANET running OSPF MANET extensions with a larger OSPF network. We first created a taxonomy of different approaches, and provided a rough bound on how many link state advertisements would be circulated in the OSPF flooding domain. We next described in detail the different approaches available, using routing techniques (OSPF hierarchy, cross-layer abstraction, BGP interconnection) and mobility management techniques (mobile mesh, readdressing). Using a mobile network emulator, we explored the performance benefit of partitioning the OSPF flooding domain via standard redistribution techniques, compared with the performance of flat OSPF/OSPF MANET routing (using a single flooding domain). We described why particular variants of OSPF redistribution are more favorable for allowing correct routing even in the face of partitions, while keeping link advertisements low when there is no partitioning. Finally, we explored the benefit of creating virtual links (tunnels) between MANET Border Routers and found the technique to perform the best in terms of minimizing the routing updates on the legacy network, while trading off some path optimality in the data plane. Detailed results of this work are found in Appendix A.
- Similarly for multicast routing, we proposed several mechanisms to connect MANET multicast to legacy networks. Specifically, we defined how MANET Border Routers need to behave to transport SMF MANET multicast traffic to a PIM-Sparse Mode (PIM-SM) region and vice versa. The design also allows multiple MBRs with little or no coordination among them. Our solution is also robust to MANET partition and changing MBR connectivity. The design allows for various optimizations. We recommend that a MANET take advantage of multiple ingress MBRs to handle network partitions quickly. We also recommend that only one node performs PIM Register messaging on behalf of its MANET, and we describe how this can easily be accomplished when using OSPF. We later extended aspects of this work to handle PIM-Dense Mode (PIM-DM) and SMF integration. Our work on PIM-SM/SMF integration is detailed in Appendix B, and PIM-DM work (unpublished to date) is presented in Section 3-5 below.

3.2 Environment

Our network simulator for this project is the Boeing Common Open Routing Environment (CORE). CORE is a user-space program that runs on a modified FreeBSD kernel that has been modified for network stack virtualization. It provides a graphical canvas that allows for drag-and-drop configuration of network topology, and animation of node mobility events. It provides a Unix shell-based execution environment for software on each node; therefore, it is a natural fit to support Boeing's OSPF MANET software or any other Unix, shell-based software (such as NRL MGEN traffic generator, tcpdump, SMF multicast routing, etc.). It is based on a fork of the open-source Imunes emulator.¹ Figure 3-1 provides a screenshot of CORE, which was developed on funding outside of the ONR RANGE contract.

¹ <http://www.tel.fer.hr/imunes/>

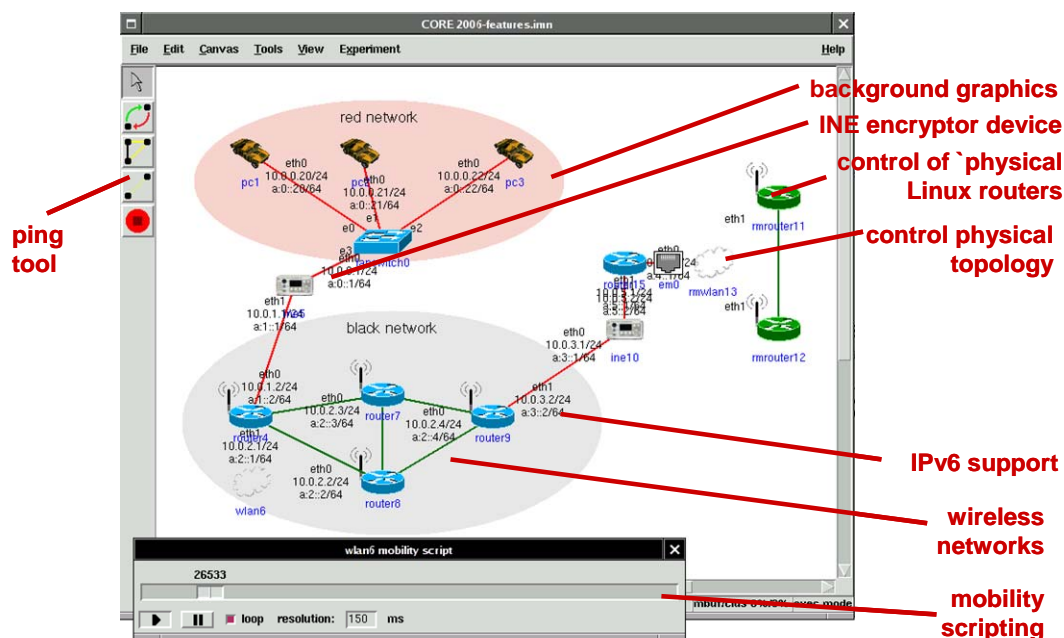


Figure 3-1: Screenshot of Boeing CORE simulation/emulation tool.

In this base program phase, we constructed an edge-based MANET topology in CORE that maps to the proposed airborne ADNS configuration for networking between wide-body aircraft, ship-based routers, and other airborne nodes. Figure 3-2 illustrates this topology. On the left is a representation of Navy ADNS shore and ship-based elements and on the right is a hypothetical multi-channel MANET. Several routers are interconnecting the MANET with the rest of ADNS; these are notionally depicted as aircraft such as E2-C, P3, and EP-3. Our emulator provides us the capability of instantiating any or all of these routers as emulated routers, or we are able to “tap” into a physical network to insert hardware devices for testing. In particular, we will insert a Cisco router into the topology at the boundary between ADNS and MANET, to explore the capability and performance offered by a standard COTS router.

Of particular interest to us with this topology is the study of how to minimize the effects of MANET routing updates on the rest of the ADNS, which has been described as a key concern of these types of network interconnections. A number of options exist and will be studied in the context of mobility and network partitioning:

- use of OSPF MANET in the airborne segment and OSPF in the ADNS segment, all within the same OSPF area;
- use of OSPF areas to segment the MANET from the ADNS;
- redistribution of OSPF MANET into OSPF; and
- interconnection of MANET with ADNS using BGP.

In particular, the first choice above does not require border router configuration but is predicted to offer the least (no) isolation between MANET OSPF (for ADNS airborne nodes) and OSPF (for ADNS shore and ships). The other techniques offer different degrees of isolation of MANET from OSPF but may not support partitioning and autoconfiguration goals as easily. We are not aware of any study that has systematically explored the problem of how to efficiently interconnect networks of these types, quantifying the tradeoffs involved. We predict the results to be immediately interesting to Navy ADNS and more generally applicable to the integration of MANET and non-MANET networks in the GIG.

We have configured the routing software to correctly handle the scenario shown in Figure 3-3. The major challenge was in dealing with the dual links that exist between the ADNS ship router and the emulated airborne node (P3 router). In FreeBSD, a longstanding known issue is that the IPv4 ARP implementation does not work well with parallel links. Our mobility scenarios were causing OSPF adjacency formation problems because the sockets directive not to consult the routing table is being overridden by the kernel, and OSPF synchronization messages are emitted from the wrong (parallel) interface. We implemented a customized kernel modification to fix this issue.

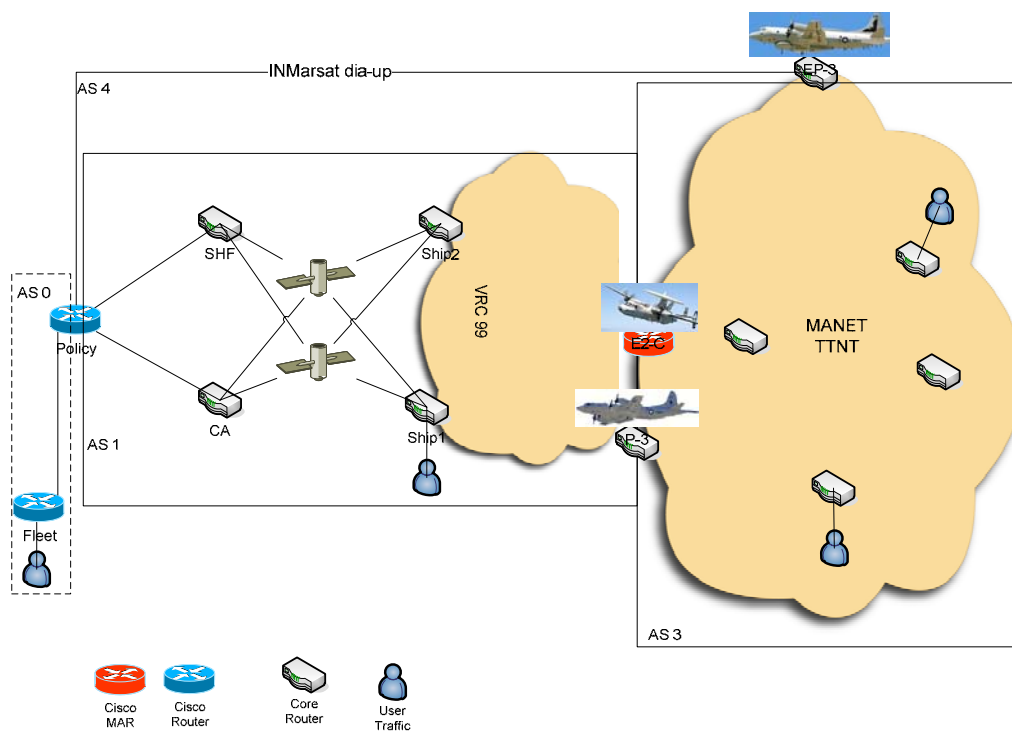


Figure 3-2: Simulation/emulation scenario of ADNS shore, ships and airborne elements for studying routing integration and route redistribution.

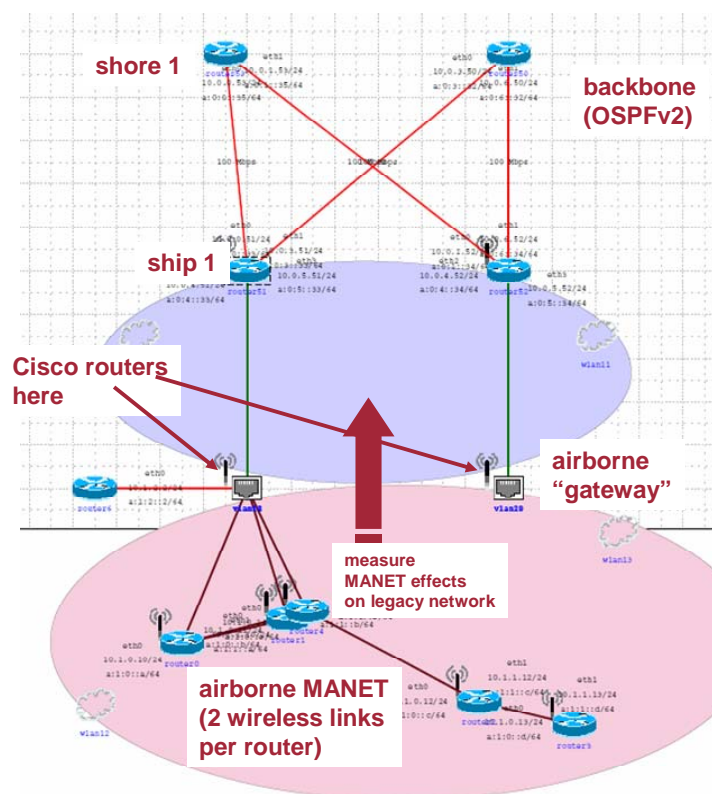


Figure 3-3: Unicast topology in the Boeing CORE environment.

We instrumented the above emulator/testbed to measure the following statistics on a per-router basis:

- neighbors/node (average node density)
- neighbor_changes/node/sec (measures network churn)
- Seconds/neighbor_life: (measures neighbor “permanence”)
- Seconds/LSA Install (measures number of LSAs that routers need to handle)
- Seconds/MANET Partition (how often a MANET partitions in some way (less than 100% MANET connectivity))

In particular, the last two statistics are of interest in quantifying the effect of mobility in the airborne segment. Partitions are a special type of mobility event in which it is necessary to convey the new routing information to the outside routers. We are not aware of other studies that have systematically explored the problem of how to efficiently interconnect networks of these types, quantifying the tradeoffs involved. We predict the results to be immediately interesting to Navy ADNS and more generally applicable to the integration of MANET and non-MANET networks in the GIG.

We have also developed a CORE-based testbed (Figure 3-4) for evaluating solutions for PIM-SM and SMF integration. The main problems that arise in multicast stem from handling multiple, independent border routers, and from incompatibilities between a group membership-based protocol (PIM) and one that does not maintain group membership (SMF). Getting multicast out of an SMF region into a PIM region requires that the multicast be sent to the rendezvous point in the PIM network so that it does not fail a reverse-path-forwarding check near the exit MANET border router. When multiple gateways are involved, there can be duplication of messages in the PIM network unless the gateways coordinate. Getting multicast into an SMF region from a PIM region can also be a challenge, since IGMP is not propagated by default in SMF. Options include flooding IGMP, tunneling IGMP to border routers, or leveraging a link state advertisement framework in the unicast protocol for group membership conveyance. Another challenge arises if packets require tagging at the border routers for SMF duplicate packet detection; if gateways are uncoordinated, the results of tagging may lead to duplicate messaging in the SMF region. We have an initial testbed running now that is (non-optimally) allowing multicast to flow into and out of the MANET, and we plan to incrementally test various optimizations to reduce duplicate messaging when multiple border routers interconnect the two regions.

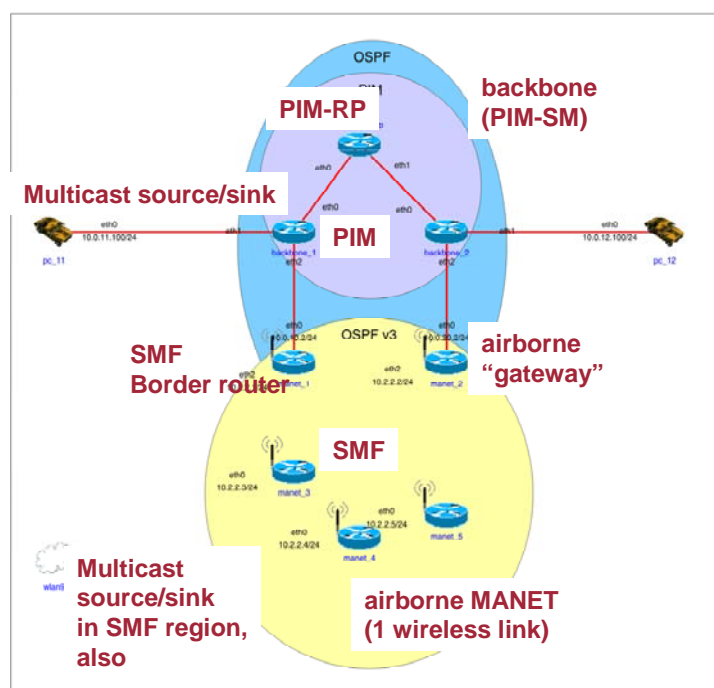


Figure 3-4. Multicast topology in the Boeing CORE environment.

3.3 OSPF MANET and OSPF Redistribution

We completed unicast data measurements related to OSPF MANET performance with OSPF route redistribution when multiple routing gateways are involved and the MANET is subject to partitioning events. The details are

summarized in the paper attached as Appendix A. Briefly, we showed in the paper that it is possible to construct a multiple-gateway tunneling approach that severely limits the amount of MANET overhead exposed to the outside network while being robust to partitioning. The below figure 3-5 shows a comparison between the overhead generated within the MANET and the overhead outside of the MANET, for a mobile MANET subject to partitioning. The solution relies on using multiple OSPF processes at the gateway to segment the flooding domains, and the use of route redistribution to summarize routes, but our findings indicated that only certain redistribution techniques were both robust to partitions and shielded the backbone from frequent routing updates.




Figure 3-5: A multi-gateway redistribution with tunnels between gateways can significantly filter the MANET-related routing overhead when propagating into an attached backbone network. The figure shows that the “Outside MANET” region experiences far fewer (higher time interval between) routing updates than the “In(side) MANET” case

3.4 PIM-SM and SMF Results

We quantified the performance of our MANET multicast gateway solution that interfaces legacy PIM-SM multicast with SMF running inside the MANET, in an emulated environment. The emulated topology is shown in Figure 3-6. In order to intuitively assess the routing path of multicast packets we set the MANET link latency to 20ms, and one of the links in the PIM network, connecting the PIM Rendezvous Point, to 100ms. A mobile node was moving inside the MANET, according to the trajectory presented in figure below, and during the mobility scenario the MANET was either merged into one component, connected through the mobile node, or partitioned into two different components.




Figure 3-6: Experimental topology depicting MANET Border Routers (MBR) interacting with Simplified Multicast Forwarding (SMF) mobile nodes and PIM-Sparse Mode routers.

We generated multicast traffic, both originated from the hosts attached to the PIM network towards the mobile node, and from the mobile node sent towards the PIM-attached hosts. Figure 3-7 shows the latency of ingress multicast traffic sent by Host 1 towards the mobile node inside the MANET. We can see that the packet delay varies with the number of hops traversed inside the MANET, as the mobile node moves, and when the network partitions and the mobile node is attached to the right component only, the latency increases by about 100 ms, as packets need to traverse the high delay link connecting the PIM Rendezvous Point. During the experiment, only one packet was lost, at sequence number 637. Uninterrupted connectivity as the network partitions is due to redundant traffic sent in our implementation to both border routers.




Figure 3-7: Ingress traffic latency – from Host 1 to Mobile.

Egress traffic follows a similar pattern, as seen in Figure 3-8. Packet delay changes with the number of wireless hops in the MANET, and also with the usage of the high latency link in the PIM network. We notice that there is a period of about 3 seconds of connectivity loss towards Host 1 as the mobile node moved from one network partition to another (around sequence number 400), due to the need for multicast tree recomputation in PIM.




Figure 3-8: Egress traffic latency –from Mobile to Host 1 (left) and to Host 2 (right).

In order to validate our protocols functionality in real networks, we instantiated a testbed composed of three Cisco routers connected to a MANET of Linux routers running our software, and demonstrated the interoperability between legacy PIM-SM multicast currently running on Cisco routers and our MANET gateway interfacing with SMF.

In the previous MBR routers, we did not run PIM and SMF on the same router. As a precursor to the PIM-DM work below, we first tested an MBR router that contained both SMF and PIM-SM. PIM-SM was used directly from the XORP distribution. Figure 3-9 below depicts the topology used. The Cisco routers ran PIM-SM with the MAR as the static RP. The Linux routers ran Xorp PIM-SM and OSPF PTMP with multicast support. These tests proved that PIM and SMF could coexist. In addition, we modified SMF to enable IGMP query and requests to be forwarded multi-hop in the MANET. This enabled PIM and IGMP interfaces on the MANET to join and leave a multicast group. To test our integration, we showed MANET ingress and egress data traffic over multiple hops. We also showed multicast sinks joining and leaving within and without the MANET.




Figure 3-9: PIM-DM and SMF testbed topology.

3.5 PIM-DM and SMF Results

We developed an initial prototype of a PIM-Dense Mode (DM) gateway that interfaces with MANET SMF and allows for multiple MANET gateways. The PIM-Dense Mode gateway does not conform to RFC 3973, but rather it offers a modified PIM-SM interface that now can negotiated with PIM-DM. The development was based on the available XORP routing software, which includes PIM-SM, but does not include Dense Mode functionality. XORP offers a very solid, widely accepted and modular platform for developing or enhancing routing protocols. The current software does not allow one to run SM and our DM interface type.

The setup used for development and testing is presented above in Figure 3-9. All nodes, other than the Cisco routers are running Linux and XORP with our software modifications. Below, we describe the current status of our DM development.

Ingress to MANET: Multicast packets sent by the Host computer (attached to the Cisco MAR router) are forwarded through the Cisco network to the MANET border routers MBR1 and MBR2, which run our PIM-DM to SMF interface. If there are subscribers inside the MANET then the MBR routers will both forward multicast data into the MANET. If Assert messages could be exchanged between the MBRs then they would agree on only one of them being the forwarder. While this is the correct behavior, in our current implementation MBR routers do not send Prune messages to upstream neighbors yet, and therefore they are unable to limit the incoming traffic within the PIM-DM network.

Egress from MANET: Packets sent from inside the MANET are forwarded through SMF to the MANET border routers (MBRs). Our PIM-DM interface running on the gateways forwards packets to the Cisco network. Our PIM interface responds and acts accordingly upon receiving Prune, Graft and Assert messages from the Cisco routers. Therefore, only one of the MBRs will forward MANET traffic, and only when there are receivers in the external network, which is the expected correct behavior.

We tested the multicast protocols on a network topology that includes a mobile network composed of four Linux routers, connected to a legacy network composed of three Cisco routers. Two of the MANET routers were acting as MANET gateways, being attached at different points to the legacy network. A Host computer was connected to the legacy network, being directly linked to one of the Cisco routers. The network topology is shown in Figure 3-10.




Figure 3-10: Network topology for PIM-SM and PIM-DM experiments.

We tested the multicast protocols in different configurations enabling one or two MANET gateways, forwarding traffic over multi-hop paths inside the MANET, disconnecting nodes and connecting them back on, and creating partitions and merges between the gateways. Routing protocols running on each box are as follows:

- Host: SMF (only for DPD)
- Cisco MAR: OSPFv2-PTMP and PIM-DM or SM
- Cisco1: OSPFv2-PTMP and PIM-DM or SM
- Cisco2: OSPFv2-PTMP and PIM-DM or SM

- GW1: OSPFv2-PTMP and PIM-SM or Boeing PIM-DM and SMF
- GW2: OSPFv2-PTMP and PIM-SM or Boeing PIM-DM and SMF
- MANET1: OSPFv2-PTMP and SMF
- MANET2: OSPFv2-PTMP and SMF

The scenarios outlined below were used to collect duplicate packet and disruption statistics for a multicast flow between a host connected to the legacy network, Host, and a mobile node, MANET1. Each experiment consisted of the following steps:

1. The source node begins sending multicast traffic at the rate of five 256 byte packets per second.
2. The destination node joins the multicast group.
3. The mobile node begins its repeated movement pattern, resting 3 minutes before moving instantly to the next position.

Four variations of each scenario were tested by changing two parameters: PIM mode, sparse mode (SM) or dense mode (DM); and multicast flow direction, into the MANET (ingress) or out of the MANET (egress). Each experiment lasted until the mobile node completed two movement cycles.




Figure 3-11: Scenario 1: Two gateways directly connected.

Scenario 1: Two gateways (GW1 and GW2) are connected to the legacy network, and directly connected with each other through the wireless channel. MANET1 moves back and forth, alternating between being directly connected to either GW1 or GW2. The experiment lasted 720 seconds and included a total of 3 moves.

Figure 3-12 shows a representative plot of multicast packets received, in this case by Host during the PIM-SM egress experiment. As expected, sequence numbers increase linearly with time, indicating that packets are received at a constant rate. Movement events occur as shown and the duration and number of packets lost during the subsequent outages were measured. Disruptions lasting less than one second were not considered outages and are not included in the results presented.




Figure 3-12: Packets received by the multicast subscriber.

Table 3-1 summarizes the overall results for each of the configurations tested. Duplicate packets only occur in the egress cases, where almost every packet is received twice. This happens because the two gateways are connected and the SMF classical flooding algorithm was used. In this case, the gateway to which MANET1 is connected receives two copies of each SMF packet: once from MANET1 and once from the other gateway. Since the current implementation does not perform duplicate packet detection before multicast forwarding is performed, two copies of the same packet are then sent into the legacy network. There is only a brief period immediately following each move in which duplicate packets are not received. This lasts until the forwarding gateway switches and an outage occurs.

PIM Mode	Multicast Direction	% Duplicate Packets	Number of Outages	Avg. / Max. Outage Duration (sec)	Avg. / Max. Outage Packet Loss (packets)
SM	Ingress	0.00	5	11.7 / 18.2	57.6 / 90
DM	Ingress	0.00	0		
SM	Egress	95.3	3	13.6 / 20.2	67.0 / 100
DM	Egress	95.3	3	20.3 / 20.4	100.7 / 101

Table 3-1: Scenario 1 results.




Figure 3-13: Scenario 2: Two gateways not directly connected in the MANET.

Scenario 2: Two gateways (GW1 and GW2) are connected to the legacy network, and not directly connected with each other through the wireless channel. MANET1 moves back and forth, alternating between being directly connected to either GW1 or GW2. The experiment lasted 720 seconds and included a total of 3 moves.

No duplicate packets are received in this scenario because the gateways are not connected and each gateway receives only one copy of each SMF packet (when connected to MANET1).

PIM Mode	Multicast Direction	% Duplicate Packets	Number of Outages	Avg. / Max. Outage Duration	Avg. / Max. Outage Packet Loss
SM	Ingress	0.00	6	9.1 / 18.8	44.7 / 93
DM	Ingress	0.00	5	10.6 / 18.4	52.2 / 91
SM	Egress	0.00	3	23.7 / 31.0	119.0 / 154
DM	Egress	0.00	3	30.4 / 31.0	151.0 / 154

Table 3-2: Scenario 2 results.

Scenario 2.5: The topology is the same as for Scenario 2, but the movement pattern is slightly different. Here MANET1 moves to an intermediate position where it is connected to both GW1 and GW2, and then moves so that it is directly connected to only one gateway (make before break). The experiment lasted 1440 seconds and included a total of 7 moves.

Duplicate packets are received in the ingress cases when MANET1 is connected to both gateways (roughly half of the time). This happens because the gateways are not connected and do not exchange Assert messages to coordinate forwarding.

PIM Mode	Multicast Direction	% Duplicate Packets	Number of Outages	Avg. / Max. Outage Duration	Avg. / Max. Outage Packet Loss
SM	Ingress	49.2	10	5.7 / 12.0	27.4 / 59
DM	Ingress	46.3	11	10.2 / 28.8	50.0 / 143

SM	Egress	0.00	3	27.1 / 31.0	135.7 / 154
DM	Egress	0.00	5	17.0 / 31.0	84.4 / 154

Table 3-3: Scenario 2.5 results.




Figure 3-14: Scenario 3 network topology, with mobile intermediate node.

Scenario 3: Two gateways (GW1 and GW2) are connected to the legacy network, and directly connected with each other through the wireless channel. GW1 is also connected to MANET2. MANET2 is connected to MANET1 over the MANET. MANET2 moves between being connected to either GW1 or GW2. The experiment lasted 720 seconds and included a total of 3 moves.

Similar to Scenario 1, duplicates packets are received for the egress cases because the gateways are connected through the MANET.

PIM Mode	Multicast Direction	% Duplicate Packets	Number of Outages	Avg. / Max. Outage Duration	Avg. / Max. Outage Packet Loss
SM	Ingress	0.00	6	8.7 / 18.4	43.7 / 91
DM	Ingress	0.00	0		
SM	Egress	95.38	3	13.5 / 20.2	66.7 / 100
DM	Egress	94.91	3	13.5 / 20.2	67.0 / 100

Table 3-4: Scenario 3 results.




Figure 3-15: Scenario 4 network topology: two gateways not connected, with mobile intermediate node.

Scenario 4: Two gateways (GW1 and GW2) are connected to the legacy network, and are not directly connected with each other through the wireless channel. GW1 is also connected to MANET2. MANET2 is connected to MANET1 over the MANET. MANET2 moves between being connected to either GW1 or GW2. The experiment lasted 720 seconds and included a total of 3 moves.

Again, no duplicate packets are received because the gateways are not connected.

PIM Mode	Multicast Direction	% Duplicate Packets	Number of Outages	Avg. / Max. Outage Duration	Avg. / Max. Outage Packet Loss
SM	Ingress	0.00	5	5.7 / 11.2	27.8 / 55
DM	Ingress	0.00	5	6.4 / 12.0	30.8 / 59
SM	Egress	0.00	3	24.5 / 31.6	121.3 / 157
DM	Egress	0.00	3	24.0 / 31.2	119.0 / 155

Table 3-5: Scenario 4 results.

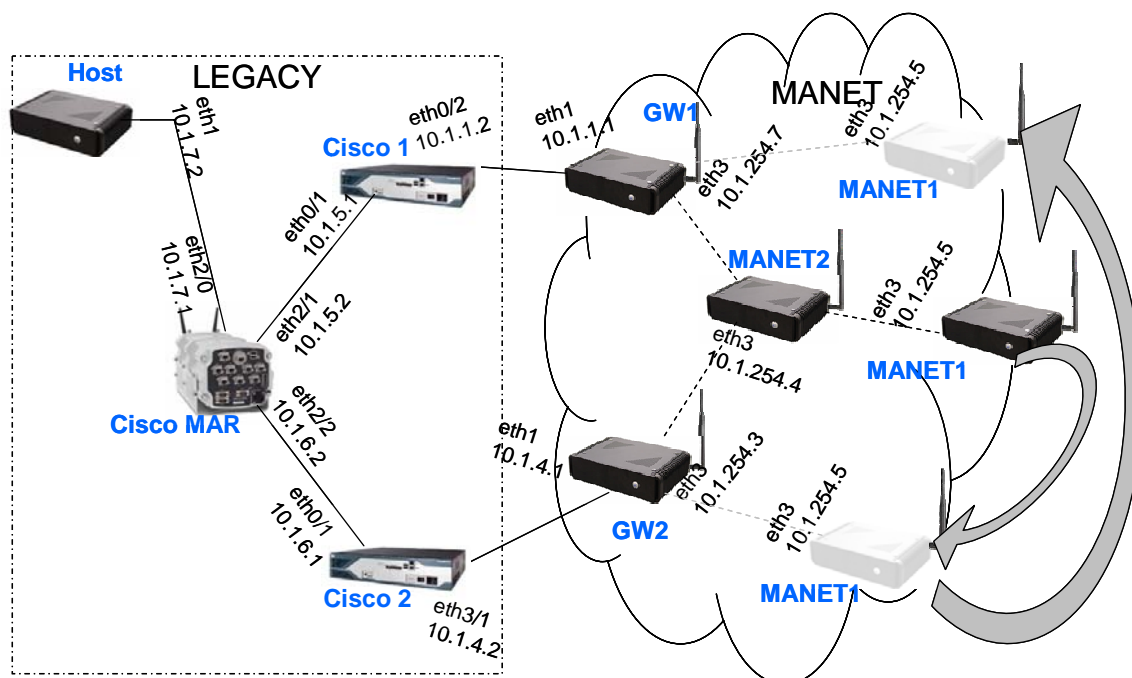


Figure 3-16: Scenario 5 network topology: Mobile moves away from intermediate gateway.

Scenario 5: Two gateways (GW1 and GW2) are connected to the legacy network, and are not directly connected with each other through the wireless channel. GW1 and GW2 are also connected to MANET2. MANET2 is connected to MANET1 over the MANET. MANET1 moves from being connected to MANET2 to being directly connected to GW2 and then directly connected to GW1. The experiment lasted 1080 seconds and included a total of 5 moves.

Duplicate packets are received for the egress cases when MANET1 is directly connected to either gateway (roughly 2/3 of the time). This happens because the forwarding gateway receives each SMF packet first from MANET1 and then from MANET2.

PIM Mode	Multicast Direction	% Duplicate Packets	Number of Outages	Avg. / Max. Outage Duration	Avg. / Max. Outage Packet Loss
SM	Ingress	0.00	12	9.4 / 30.4	45.8 / 151
DM	Ingress	0.00	6	7.8 / 15.4	38.3 / 76
SM	Egress	63.0	6	12.0 / 20.2	60.0 / 100
DM	Egress	63.0	5	14.2 / 20.4	70.4 / 101

Table 3-6: Scenario 5 results.

4 Improved Routing Software

4.1 Summary

In this program, we developed unicast and multicast routing extensions to existing routing software packages. All of our routing software is provided as open source. We have been extending existing routing suites (quagga, XORP, and NRL SMF) and our modifications are available under the existing open source license terms as derivative works.

- We continued to develop OSPF-MANET code based on the quagga OSPFv3 implementation. We developed an Address Families extension to allow route redistribution between OSPF-MANET and quagga OSPFv2.
- We continued work on getting SMF to interwork with XORP-based PIM implementation, and on IGMP propagation through an SMF-based MANET. We developed a diverter agent at the SMF border router to allow multicast to flow in and out of the MANET. We modified the PIM-SM implementation to interface with a Cisco PIM-DM router upstream, to conduct PIM-DM/SMF testing.

4.2 Unicast Routing Software: OSPF MANET

Our OSPF-MANET implementation began under a previous contract effort (ONR KSA FNC Block 2). OSPF-MANET is a set of extensions to the OSPFv3 implementation of the quagga routing suite. To maintain consistency with Quagga Routing Suite, the quagga OSPF-MANET code has now been ported to the most recent stable version, 0.99.9. We also performed the following work:

- We released our OSPF-MANET code integrated with SMF multicast to NRL. Prior to release of the code, we tested the integration of OSPF-MANET-SMF in CORE. SMF was shown to forward multicast packets using the MDR set generated by OSPF-MANET. We also provided our CORE emulation tool (not developed under government funding) to NRL in form of a DVD containing a VMware image that should be able to be easily run in Windows.
- We assisted Jeff Weston (NRL) in getting the new OSPF-MANET code building for NRL since he ran into some problems with building quagga OSPF on a Fedora Core 6 box. We worked with him to enable him to build and run quagga on Fedora Core 6.
- We made a number of minor improvements to the OSPFv3 MANET implementation, including fixing some logging bugs, additional logging for data collection, removing leaf nodes from the forwarding set advertised to SMF, adding support for parallel links between routers in OSPFv3 with AF, and fixing a bug that did not add the loopback interface as the outgoing interface when the destination was directly connected in OSPFv3 with AF.

On the route redistribution front, we have added support for advertising varying type 1 LSA costs external to an AS. Next, we improved the interface between OSPF-MANET and SMF to use the built in “pipeExample” to create a cleaner and more portable interface. Finally, we fixed various bugs due to operating systems specific abnormalities in FreeBSD and Fedora Core Linux.

4.3 Multicast Routing Software: SMF and PIM

SMF has been developed and maintained by NRL, and PIM-Sparse Mode (PIM-SM) is available under the XORP routing suite. We developed a PIM-Dense Mode (PIM-DM) interface by modifying the XORP PIM-SM implementation, and we also performed software work oriented towards getting SMF and XORP PIM to interwork in a gateway configuration.

- We worked on getting SMF to interwork with PIM on the mobile router, and on IGMP propagation through an SMF-based MANET. We developed a diverter agent at the SMF border router to allow multicast to flow in and out of the MANET, although as of the beginning of April we have not completed the IGMP propagation piece that allows a MANET router to signal join requests to the gateway, and static forwarding is required in that direction.
- We also focused on getting SMF to interwork with PIM on the mobile router, and on IGMP propagation through an SMF-based MANET. We completed the basic PIM/SMF multicast integration scenario in CORE for RANGE. We completed an initial multicast scenario for our CORE emulator as well. The CORE emulation consists of two PIM routers, three MANET routers (MRs), and one MANET border router (MBR). The MRs flood IGMP report messages to the MBR. The MBR delivers these IGMP reports to the attached PIM router to inform the PIM router of the MANET's multicast membership. The MBR is also

responsible for injecting MANET multicast traffic into the PIM area, and non-MANET multicast into the MANET. This scenario demonstrates the current best practice and how fragile they are. We created internal documentation describing the multicast scenario, tools, and demo.

- We developed an initial prototype of a PIM-Dense Mode (DM) gateway that interfaces with MANET SMF and allows for multiple MANET gateways. The PIM-Dense Mode gateway does not conform to RFC 3973, but rather it offers a modified PIM-SM interface that now can negotiate with PIM-DM. The development was based on the available XORP routing software, which includes PIM-SM, but does not include Dense Mode functionality. XORP offers a very solid, widely accepted and modular platform for developing or enhancing routing protocols. The current software does not allow one to run SM and our DM interface type.

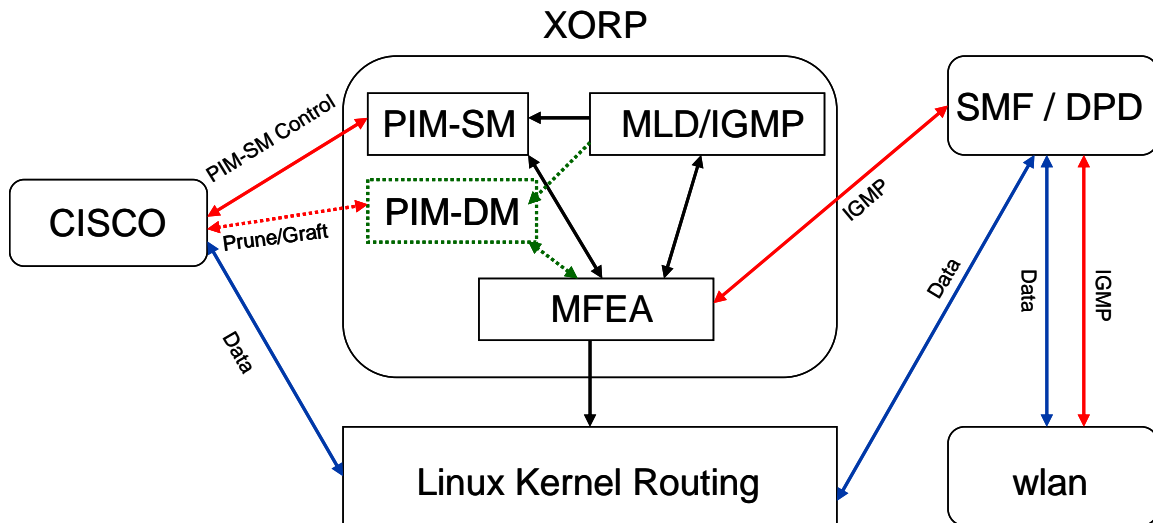


Figure 4-1: Functional description of how SMF and XORP coexist at a multicast gateway

5. Option Planning

5.1 Summary

During the base program phase, Boeing explored a number of options for conducting field testing and demonstrations of RANGE software in the option phase. Boeing settled on the University of Illinois at Urbana-Champaign as the most cost-effective option for this program, and plans to conduct a demonstration with two UAVs, augmented by ground and emulated nodes, in the summer of 2008. Boeing also developed, with NRL, concepts for additional D&I research goals for the option phase. This option was exercised by ONR on January X, 2008.

5.2 Field Demonstration preparation

Boeing described three options for UAV internetworking demonstrations;

- 1) use 2 real UAVs (supported from University of Illinois at Urbana-Champaign), some additional emulated UAVs and multiple ground nodes, with total estimated cost controlled (\$150K max);
- 2) continue to explore piggyback opportunity of other ONR program UAV demonstrations, for example, Deep Lightning Bolt (DLB) Project B - AIS on ScanEagle UAV program;
- 3) consider a fallback option of Lab-based UAV demonstration in an emulated environment focusing more on scientific experiment and test.

Santanu Das of ONR voiced support for option 1, asked to use more emulated UAV nodes as necessary, and to add if possible, other type of radios (e.g., VRC-99, TCDL) than 802.11a/b/g radio for interoperability of heterogeneous link technologies. Chris Rigano agreed and also mentioned potential issues of UAV EMI and frequency coordination. He also mentioned the possibility of using a government-furnished emulation tool.

Boeing explored whether Boeing/Insitu ScanEagle might be available for demonstration. Boeing received an estimate from Insitu on the cost of a ScanEagle demonstration based on the SOW, but it exceeded the amount allotted for this demonstration.

Therefore, Boeing completed subcontract negotiations with UIUC regarding the use of Prof. Natasha Neogi's (neogi@uiuc.edu) UAVs in a field demonstration to be conducted at a UIUC site in 2008. Below is a figure of one of the two UAVs to be used in the demonstration. The UAV (Figure 5-1) has approximately a 78" wingspan with a payload capacity of roughly 5 lb, and a flight time of 15 minutes currently (30-45 minutes projected). The UAVs use a 900Mhz radio for Piccolo autopilot, embedded a PC-104 computer (ULV Celeron 400Mhz), and have a video capability. The plan is to equip these UAVs with Boeing routers running RANGE software, and to augment the two UAVs with additional physical and emulated nodes on the ground.




Figure 5-1: UIUC UAV.

Figure 5-2: UIUC UAV set for field testing of routing protocols

In July 2007, Boeing visited UIUC, where we discussed with Prof. Natasha Neogi the UAV capabilities, integration of our routing protocols within the UAV hardware, and the field testing plan. The UAV allows manual remote control, and is also capable of autonomous operation, being driven by a Piccolo autopilot. Currently, the UAV has an on-board D-Link router that was used for simple communication testing over 802.11 wirelesses. The pictures presented in Figure 5-2 show the UAV on the ground, during field testing, and also the internal placement of radio components and payload bay for our single-board computer.

We tested the UAV in flight, and verified ground to air 802.11g performance. Using a regular laptop computer pinging the D-Link router installed on the UAV, the channel exhibited relatively low signal to noise ratio (Windows XP reported about 10% quality); however, connectivity was maintained for the duration of a complete flight, with a packet loss rate of about 10%. We believe that better performance can be achieved by lowering the transmission rate and by increasing the transmitting power of the wireless devices. In the planned field testing, we will replace the D-Link router with a Linux based router running our suite of mobile protocols.

5.3 Future D&I Topics

We have developed the following preliminary plan for D&I work in this option. We plan to devote roughly 2/3 to 3/4 of the effort towards multicast networking, since we perceive that it is a more critical need for RANGE scenarios. We plan to focus on two topics:

- Unicast routing over low-data-rate bearers
- Policy-based interconnection of multicast regions

5.3.1 Unicast routing

Most research work on OSPF MANET has been conducted in the context of IEEE 802.11-based ad hoc radio networks. While IEEE 802.11-based radios (commercial and tactical) are likely to be a predominant radio type for OSPF MANET, there is interest in understanding how well it could help low data rate radio channels such as UHF and HF systems. For instance, experiments in coalition networking are using OSPF as the heterogeneous routing protocol over different radio links, and experiments continue on using HF-IP and UHF (Subnet Relays) as bearers for these networks. OSPF MANET, or modifications to OSPF MANET, may improve the number of nodes that can be accommodated on these channels.

An important consideration in such networks is how large numbers of external link state advertisements (LSAs) may be efficiently handled. Because OSPF is vulnerable to scalability problems for very large routing domains, the operational tendency is to segment the network into smaller routing (flooding) domains and redistribute routes between the domains. This causes a large number of external LSAs to be circulated over each network; often more LSAs than are generated internally. Therefore, this study should consider the scaling effects of external LSAs in the system, and ways to mitigate large numbers of them. Within scope of this study are possible protocol adjustments (tuning), protocol changes (e.g., compression), and other novel techniques (e.g., different handling of internal vs. external LSAs) to improve performance in low bandwidth scenarios.

Finally, we will perform one or more case studies (e.g. SNR, HF-IP) of practical use to the Navy, considering the specific characteristics of deployed links. The value to the Navy of this research is potentially improved performance of OSPF MANET software and protocol in low bandwidth scenarios, and estimates of scalability of future deployments of UHF or HF-based networks, contrasting OSPF MANET performance with non-optimized legacy OSPF performance

5.3.2 Policy-based interconnection of multicast regions

Interconnection of multicast routing regions is also important, based on the RANGE networking CONOPS, but it has received relatively little research emphasis due to lack of commercial deployment of multicast. In the base program phase, we emphasized the raw interconnection of SMF and PIM routing domains, but we did not provide any way to control or filter multicast datagrams between the two domains. Controlling the flow of multicast data between routing regions is important for performance reasons but also for policy reasons. In the option phase, we plan to study aspects of policy-based interconnection of multicast domains, including coordination between multiple gateways that may become partitioned from one another.

Because policy-based routing interconnection is a very large topic, we plan to implement and study incremental capabilities for policy-based control of multicast interconnection, including the following:

- **Policy capability 1:** Provide group membership to the MANET border via multicast OSPF (MOSPF)-like technique (involves building in an opaque LSA to OSPF and a client interface to the quagga process to set/get these LSAs). Enable/disable flooding based on presence of group members in a particular domain.

- **Policy capability 2:** Allow per-group configuration for either "ALL_MANET_GATEWAYS" flood or "SINGLE_MANET_GATEWAY" flood. SMF will read in a flat policy configuration file and, if there are group members, will either elect a single gateway to flood, or will use all reachable gateways. Requires inter-gateway messaging protocol (likely based on opaque LSA types). In this phase, the gateway elected by the SINGLE_MANET_GATEWAY is not based on traffic levels or distance from group members to the respective gateways (e.g., each gateway knows which are all the other gateways, and the lowest lexicographic gateway, based on a hash of the group address and other things, enables forwarding)
- **Policy capability 3:** Allow SINGLE_MANET_GATEWAY election above to be influenced (weighted) by presence of nearby receivers (learned through the capability #1 above)

We plan to study and show some tangible benefits of these policy capabilities (illustrated below in Figure 5-3) with different applications (e.g., some applications may benefit from the robustness of ALL_MANET_GATEWAY, while others may be fine with SINGLE_MANET_GATEWAY. We will also study whether these extensions save bandwidth, and if so, how? Along the way, we will further develop the multicast (and combined) software in these areas:

- OSPF Opaque LSA mechanism
- OSPF/SMF interface improvements
- XORP/SMF/DPD interaction
- Better PIM-DM implementation

The value to the Navy is that multicast seems likely to be as prevalent or more so than unicast for GIG tactical edge applications, and the lack of multicast software has been cited as a barrier to deploying and testing tactical edge experiments.




Figure 5-2: Policy-based interconnection of multicast routing domains

References

- [ADN05] "Airborne ADNS (aADNS) Routing Working Paper," ADNS Engineering Lab, October 18, 2005.
- [Cha07] I.D. Chakeres, C. Danilov, T.R. Henderson, and J.P. Macker, "Connecting MANET Multicast", IEEE MILCOM Conference, October 2007.
- [DeepLightningBolt] Deep Lightning Bolt Project B, 2006.
- [LA05] LT Sumner Lee and George Arthur, "Airborne ADNS," briefing for ADNS, December 2005.
- [MDC04] J. Macker, J. Dean, and W. Chao, "Simplified multicast forwarding in mobile ad hoc networks," in IEEE MILCOM Conference, 2004.
- [Ogi07] R. Ogier and P. Spagnolo, "MANET Extension of OSPF using CDS Flooding," Internet-Draft (work in progress): draft-ogier-manet-ospf-extension-10, IETF, November 2007
- [RFC2328] "OSPF Version 2," Internet Engineering Task Force (IETF), Request for Comments: 2328, April 1998.
- [RFC2362] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei, "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification," RFC 2362, June 1998.
- [RFC3793] A. Adams, J. Nicholas, and W. Siadak, "Protocol Independent Multicast - Dense Mode (PIM-DM), Protocol Specification (Revised), RFC 3973, November 2005.
- [Spa07] P. Spagnolo and T. Henderson, "Connecting OSPF MANET to Larger Networks", IEEE MILCOM Conference, October 2007.
- [TEN06] "Global Information Grid Joint Tactical Edge Networks (JTEN) Engineering White Paper," Draft Version 1.3, 10 July 2006.

CONNECTING OSPF MANET TO LARGER NETWORKS

Phillip A. Spagnolo and Thomas R. Henderson

Boeing Phantom Works
P.O. Box 3707, MC 7L-49
Seattle, WA 98124-2207

Abstract— The Open Shortest Path First (OSPF) routing protocol for IP networks performs inefficiently when operated over multihop, broadcast-based radio channels such as IEEE 802.11 in IBSS mode. These types of networks are known more generally as mobile ad-hoc networks (MANET). The Internet Engineering Task Force (IETF) OSPF working group has been considering extensions to OSPF to enhance performance over such channels, and multiple OSPF MANET proposed extensions are advancing to experimental status. Although OSPF MANET performs more efficiently within the MANET itself, it has the side effect of propagating large numbers of routing updates throughout the rest of the OSPF flooding domain. In this paper, we examine OSPF MANET's interaction with attached non-MANET networks, and explore various techniques for reducing the overhead, such as static and dynamic route summarization, route redistribution, cross-layer abstraction techniques, and tunneling approaches between MANET Border Routers (MBR)s. We use a hybrid testbed/emulated network to quantitatively explore the tradeoffs between different approaches.

I. INTRODUCTION

Mobile Ad Hoc Networks (MANETs) are a class of computer networks characterized by the use of wireless communications and the presence of dynamically changing connectivity [1]. MANETs can operate in the absence of fixed infrastructure, but often are deployed as extensions to more stable network topologies.

Extensions for MANET routing to the Open Shortest Path First (OSPF), [2] and [3], routing protocol were evaluated in [4]. The paper explained that extending OSPF for MANET would allow MANET routing principles to be applied to enterprise routing protocols, potentially easing the operations and management burden compared with operating separate protocols. Two extensions, MANET Designated Routers, [5] and Overlapping Relays, [6], were implemented, simulated, and examined. Both extensions used three basic strategies to reduce routing overhead: flooding, adjacency, and topology reduction, and both made significant strides to integrate enterprise and MANET routing.

Now that MANETs can be better integrated with enterprise networks, a problem has become more acute. Most work on MANET routing has focused on reducing overhead and making routing robust within the MANET, but for the most part, this work has ignored implications on the global network due to the frequently changing MANET. This work explores a variety of methods to suppress the frequently changing link state information within the MANET from being exposed to the larger network while maintaining accurate and efficient paths into the MANET.




Fig. 1. OSPF flooding from a notional airborne network through MANET Border Routers into the larger network.

This paper is organized as follows. First, we will discuss the repercussions of linking a MANET with a legacy network. Next, we will outline a number of approaches to limit the effect the MANET has on the larger network. Then we will use emulation and implementation to validate our conclusions. Finally, we will summarize the research performed in this paper.

II. IMPACTS OF OSPF MANET ON LARGER OSPF NETWORKS

Figure 1 provides a notional network diagram of an airborne, highly mobile OSPF MANET connected to a larger slowly changing OSPF network. When the MANET topology changes, link state advertisements (LSAs) will be flooded by OSPF throughout the network to maintain accurate routes. The rate at which the link state information is flooded is dependent upon the rate at which links change. The link change rate varies based on network density and mobility.

In section III.D.1 of [4], we described that the rate at which LSAs were originated bumped up against OSPF's architectural maximum for LSA generation, minLSInterval , in many mobility scenarios. This means that a new LSA was generated every minLSInterval seconds by each router in the network. OSPF's default minLSInterval is 5 seconds. If the OSPF network is flat (no hierarchy), then the maximum rate at which a new LSA will be flooded into the larger network is $\frac{\text{NumOfRouters}}{\text{minLSInterval}}$. Figure 2 and Equation 1 depict the average number of seconds between LSA origination by a router in a MANET as a function of the neighbor changes per node per second. Note that the limits are the LSRefreshTime for zero mobility and MinLSInterval for high mobility. The variable 'a' in the equation is a constant that is influenced by the distribution of neighbor changes. If the neighbor changes




Fig. 2. Average Number of Seconds between LSA origination in a MANET.

occurred periodically every minLSInterval then 'a' would be equal to one.

$$LSAOriginationTime = \begin{cases} 1800sec, & x = 0 \\ \frac{1}{ax}, & 0 > x \geq 0.2 \\ 5sec, & x \geq 0.2 \end{cases} \quad (1)$$

A stated goal of OSPF MANET is to allow vendors to provide a MANET-capable routing protocol within the well-established OSPF protocol and network management framework. However, if the cost of supporting efficient operation in the MANET results in the rest of the OSPF domain being inundated with routing updates, the benefits may be all for naught. This observation leads to the question of whether mechanisms are available to limit the impact of MANET routing changes on the larger network, without seriously impacting performance.

The first logical idea is to simply configure the OSPF-MANET as an area and statically summarize the network's address prefix(es). This works very well when there is a single MANET Border Router (MBR) to the MANET because routing follows a single path into and out of the stub area; thus LSAs do not need to be advertised for every link change in the MANET.

However, Figure 3 depicts a MANET with two MBRs to the larger network. Summarization works fine in this case too until the network partitions, then the legacy network does not have enough information to choose the correct MBR to enter the MANET because its knowledge of the MANET has been summarized. Therefore, OSPF routes may lead to forwarding packets into the wrong portion of the MANET. Furthermore, summarization may still be possible if it is done dynamically; however, it is unlikely that address prefixes will partition along convenient boundaries so that summarizing prefixes is even possible.

III. APPROACHES TO LINK OSPF MANET AND NON-MANET NETWORKS

Using OSPF-MANET within a larger OSPF network presents two competing design goals from the perspective of the larger OSPF network:

- 1) To create routes that take the minimum cost path from a source outside the MANET to a destination in the




Fig. 3. Partitioned network topology with a MANET Border Router in each side of the partition.

MANET or the minimum cost path to traverse a transit MANET.

- 2) To minimize the amount of link state information originated within the MANET due to frequent link changes that is flooded into the larger network.

Put simply, there is a tradeoff between how much abstraction can be performed to minimize the visibility of the MANET mobility and how well path selection performs. The larger network can be impaired if the frequently changing MANET generates too much overhead, but the MANET can be rendered usable if data is taking unnecessary long paths. Unnecessarily long paths in a MANET can be considered part of the total routing protocol overhead [7].

We described above in Section II that in the simple case of a MANET stub network with a single MBR to the larger OSPF network, it is sufficient to summarize the attached MANET behind a stable advertised address prefix, and allow the MANET nodes to find a default route to the router. In this case, there is not really any OSPF-specific nature to the problem; any unicast MANET routing protocol can be integrated in this way. The challenge lies in the case when there are multiple MBRs in a MANET and there exist multiple ingress and egress points to the MANET. The solution must be able to find the right balance between abstracting mobility-induced events from affecting the OSPF core, while also providing enough robustness to select the correct MBR entry point to handle partitions and lowest cost routing.

In this section, we provide a brief taxonomy of available approaches to this problem, and in the next section, we provide some quantitative evaluation of the use of OSPF mechanisms alone to solve this problem. The following approaches are divided into two classes of techniques: routing and mobility management.

A. Routing Techniques

1) *OSPF Hierarchy*: The OSPF protocol allows a network to be separated into areas or autonomous systems (AS)es, which we will define here as OSPF "regions" and the routers that interconnect them as OSPF border routers. A defined set of OSPF procedures governs how routes may be advertised and used between regions. The procedures allow OSPF regions to

be summarized by border routers. The routes advertised are called redistributed routes.

There are three techniques used to redistribute routes from one OSPF region to another. We will classify them here by LSA Type. Type 3 LSAs, otherwise known as Summary LSAs, allow routes to be exchanged from one OSPF area to another. Type 5 LSAs, known as External LSAs, allow routes to be exchanged from one OSPF autonomous system to another. Of these External LSAs, there are two variants: External Type 1 LSAs allow a router to assume that the advertised metric may be added to metrics in use internally within the AS. External Type 2 LSAs must be handled differently by OSPF routers computing shortest paths: metrics within the AS must not be added to External Type 2 LSAs, and higher path preference is always given to paths not using External Type 2 routes.

Finally, OSPF border routers can be configured to advertise fixed or dynamic cost metrics on redistributed routes. Fixed costs mean that all routes from other regions are advertised as being the same cost from the border routers, and dynamic costs allow the actual costs to be advertised.

With the above introduction, we first note that static OSPF route summarization and redistribution of the MANET will not achieve the goal of surviving network partitions, because during a partition, OSPF border routers will advertise reachability to the aggregate but not be able to actually reach all MANET destinations. This leaves six routing options to redistribute routes without summarization.

Table I lists six possible approaches to OSPF route redistribution that do survive partitions. The different approaches trade off routing accuracy for reduced advertisement overhead. In the table, column two gives the LSA type used by the border router, and column three states fixed or dynamic metric. Column four indicates the path chosen to enter the MANET. Column five shows what type of MANET topology changes are exposed outside of the MANET. Exposing partitions is the minimal requirement for the redistribution scheme to survive partitioning.

In each of the six options, an LSA is originated for each route within the MANET. For OSPF-MANET this is on the order of the number of nodes times the number of advertised prefixes per node. Three of the six options, 1, 3, and 5, expose all changes in the MANET to the larger network. This would constitute LSA flooding according to Equation 2 and would provide no improvement over flat routing. Of the remaining three, only options 2 and 6 expose partitions and give the least cost path to the MANET. Therefore, they give the best tradeoff of the design goals. Option 2 can be used for redistribution between OSPF ASes, and Option 6 can be used for redistribution between OSPF areas. They both provide the shortest path to get to the MBRs, and each only propagates LSAs in the larger network when a route to a prefix in the MANET is lost or gained. This is because the route to routers in the MANET may change paths inside the MANET, but the cost to reach the destination is fixed in the advertisement. It is important to note that certain implementations of OSPF send out an LSA even when it is not necessary in this case. Cisco routers do not originate an LSA when the cost is static.

The penalty to maintain this property is that suboptimal path selection into the MANET is performed because visibility into the MANET from the larger network is blocked. In Section IV, we quantify the performance changes due to these techniques.

RC = Route change to a prefix

R = Route added or deleted to a prefix

P = Number of prefixes advertised by the router

$$\frac{\text{FloodedLSAs}}{\text{sec}} = \left(\frac{RC}{\text{sec}} \right) \left(\sum_{\text{MANETRouter}0}^{\text{MANETRouter}N} P \right) \quad (2)$$

$$\frac{\text{FloodedLSAs}}{\text{sec}} = \left(\frac{R}{\text{sec}} \right) \left(\sum_{\text{MANETRouter}0}^{\text{MANETRouter}N} P \right) \quad (3)$$

Opt	LSA Type	Metric	Gateway Entry	Exposes
1	Extern Type 1	dynamic	best end-to-end route	All Changes
2	Extern Type 1	fixed	lowest cost to GW	Partitions
3	Extern Type 2	dynamic	lowest cost from GW	All changes
4	Extern Type 2	fixed	any connected GW	Partitions
5	Summary	dynamic	best end-to-end route	All changes
6	Summary	fixed	best MANET entry	Partitions

TABLE I
OSPF HIERARCHICAL ROUTING TECHNIQUES

2) *Cross-layer Abstractions*: The previous section described ways to reuse the hierarchy and redistribution techniques available in the OSPF protocol, to hide selected mobility events from the rest of the network. Another approach is to use cross-layer integration techniques to abstract the mobility to a lower protocol layer. Consider an approach in which OSPF is implemented on top of a layer-2 protocol that implements MANET routing/bridging. In such an approach, the system can be designed, and the protocol layers coordinated, such that the layer-2 protocol handles the mobile routing events, and works to produce the appearance of a completely-bridged network to layer-3 (i.e., in the simplest case all nodes at layer-3 appear to be one-hop away from one another, while in reality they may be multiple layer-2 hops away from one another). Such an arrangement, if coordinated correctly between the protocol layers, can allow the layer-3 topology (advertised to the rest of the OSPF network) to change less frequently even though the underlying layer-2 topology may undergo more rapid change. A well-known example of this type of design is the Radio OSPF (ROSPF) protocol implemented on top of a mobile Intranet layer [8].

Although cross-layer techniques have the potential to provide superior performance, they are not available for all radio technologies due to interlayer dependencies. Hence we have focused in this paper on solutions at the IP layer and above.

3) *BGP Interconnection*: The Border Gateway Protocol (BGP) allows for a policy-based interconnection of Autonomous Systems. BGP is a path-vector protocol [9], which means that the routing tables maintain the paths of ASes traversed to reach the end systems, and the protocol allows for a number of techniques to control and filter the paths

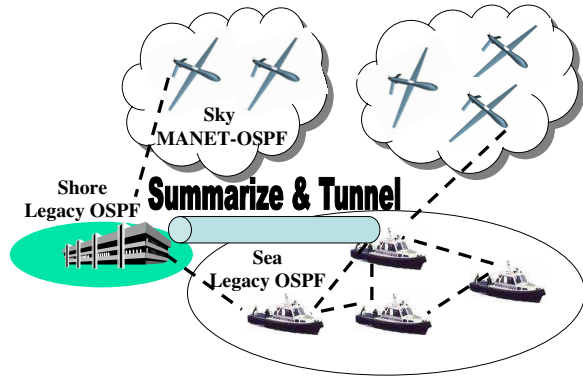


Fig. 4. Summarization at MBRs and tunneling between MBRs

to achieve policy-based objectives. The use of BGP to interconnect a MANET with a legacy OSPF system is possible, and has the benefit of shielding the OSPF network from the MANET updates. However, BGP was not originally optimized for rapid convergence in the face of numerous routing updates. In fact, BGP builds in a number of mechanisms to dampen the impact of routing updates on the system. There are also questions about the scalability of a large internal BGP network to interconnect many MANETs. The net result is that we consider BGP to be a less promising protocol for MANET interconnection than other techniques discussed herein.

B. Mobility Management Techniques

It is possible to solve the multiple MBR problem outside of pure routing solutions. We briefly mention two such techniques.

1) *Mobile Mesh and Tunneling*: If it is possible for the MANET border routers to communicate with each other over the non-MANET links, then virtual links (tunnels) can be created between them over the non-MANET links (see Figure 4). This allows each MANET border router to summarize the MANET prefixes and advertise them to the larger network because a packet that reaches any MBR can be tunneled to another if needed. An MBR partitioned from a MANET router survives the partition by tunneling across the non-MANET links to another MBR that has connectivity to the router. The tunneling method also enables packets to enter the MANET from the MBR with the lowest cost route to the destination. An MBR that does not have the lowest cost route to the destination forwards packets to the lowest cost MBR. The penalty for using tunnels is the overhead to establish and maintain the tunnels, and the additional delay and bandwidth usage over the tunnel. A suite of protocols and a research implementation, known as “Mobile Mesh,” has been developed along these lines [10]. This type of solution has the potential of shielding the legacy OSPF network from all mobility induced events, even partitions, so long as the border routers themselves do not become partitioned with respect to the legacy network.

2) *Readdressing*: Finally, approaches are possible in which mobile nodes affiliate with a MANET border router that serves also as a DHCP server or IPv6 router issuing Router Advertisements. When reachability to the affiliated border router is lost, affected nodes could reestablish connectivity to a

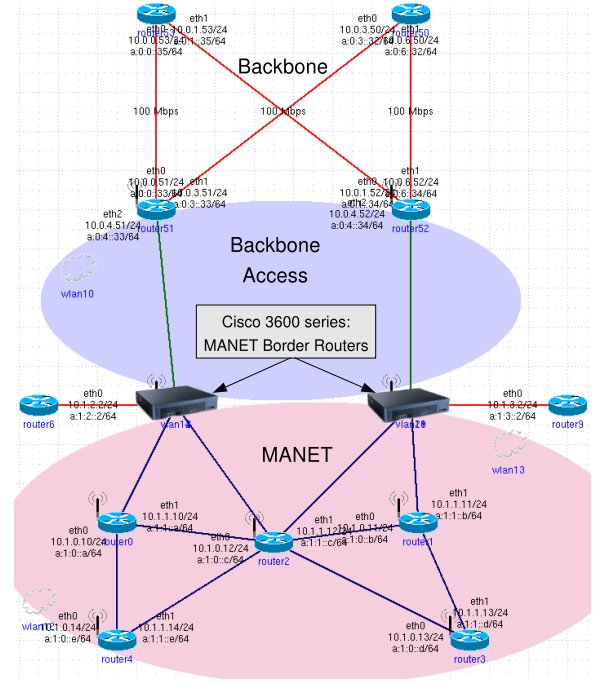


Fig. 5. CORE emulation scenario

new border router and readdress their MANET interface within the new border router’s prefix. If mobile hosts are supported in such a manner, some means to migrate transport connections to new addresses (such as Host Identity Protocol mobility [11]) is needed, while if mobile routers changed network reachability and readdressed the router interface, some type of mobile router [12] capability could be used.

IV. NUMERICAL RESULTS OF OSPF REDISTRIBUTION TECHNIQUES

Boeing’s Common Overlay and Routing Environment (CORE), a variant of the IMUNES emulator [13], was used to quantitatively evaluate techniques to reduce the effect of rapidly changing MANETs in OSPF. CORE allows interconnection of (real) Cisco and (virtual) Quagga routers in a real-time emulated network. See Figure 5 for a screen shot of CORE and the emulation scenario. The circular light blue routers are OSPF Quagga routers, and the rectangular dark blue routers are Cisco 3600s running OSPF. In each scenario, a backbone legacy network is linked with a MANET through a backbone access network. The backbone network emulated wired links, the backbone access network emulated wireless stable links, and the MANET emulated parallel frequently changing wireless links. OSPF broadcast interfaces were used on the wired links and point-to-multipoint links were used on wireless links. The MANET was scripted within CORE to move at varying change rates, to emulate a highly mobile network connected to a stable larger network.

We first provide results for a baseline scenario in which a single OSPF area contains the larger network and the MANET. Next, we use redistribution between OSPF ASes to limit the MANET impact on the backbone. Then, we use redistribution with tunneling between MBRs to enable summarization of the MANET prefixes. Each data point, in figures 6 through

8 below, was obtained by running the scenario for over two hours. The scenarios were run multiple times varying the wireless range on the MANET links after each run. Unless otherwise specified, the parallel wireless links have the same range. The following statistics were collected:

- 1) *Average Neighbors / MANET Node*: The average number of neighbors of each Quagga MANET router, computed by tracking the proportion of emulation time that each router has a given number of neighbors. This statistic measures the average network density.
- 2) *Seconds / MANET Neighbor Change / MANET node*: The average number of seconds between neighbor changes on the MANET Quagga routers. This measures the network churn.
- 3) *Seconds / MANET Neighbor Lifetime*: The average number of seconds that a Quagga MANET router's neighbor is in state two way or greater. This statistic measures neighbor permanence.
- 4) *Seconds / MANET LSA Install In MANET*: The average number of seconds between installing a MANET-originated LSA, measured at a router within the MANET.
- 5) *Seconds / MANET LSA Install Outside MANET*: The average number of seconds between installing a MANET-originated LSA, measured at a router outside the MANET. This statistic measures the LSA flooding impact of the MANET on the larger network.
- 6) *Seconds / MANET Connectivity Change*: The average number of seconds between losing or gaining (OSPF-measured) connectivity to a router in the MANET, as measured at a border router (MBR). This statistic measures how often the MANET partitions and merges.

In the first scenario (our baseline), a MANET is connected to a larger OSPF network. The MANET and the larger network are in the same area and no special configuration is performed on the MBRs. In this case, the MANET is completely exposed to the larger network and vice versa. The baseline scenario maximizes our first design goal of low cost routing through the MANET, but it exposes all link state information outside the MANET which is counter to our second design goal. Figure 6 shows a plot of the average number of seconds between a router installing an LSA that advertises the MANET. In this figure, there are two overlapping lines that represent a router installing an LSA inside the MANET and outside the MANET. The overlapping lines show that the MANET changes affect the inside and outside of the MANET equally. In this baseline scenario, an LSA from the MANET is installed every 1.5 seconds on average. This interval between advertisement installs outside the MANET is what we desire to increase, without compromising performance.

In the second scenario, the larger network and the MANET are placed in two separate OSPF ASes. The MBRs (Cisco routers) perform redistribution of External Type 1 LSAs with fixed cost to limit the effect of MANET LSA flooding on the larger network. In this setup, LSAs will only be flooded between the two ASes when a partition or merge occurs in the MANET. One technical note must be made about configuring




Fig. 6. Larger network linked with a MANET in a single area OSPF

Cisco OSPF routers. The administrative distance of routes in the routing tables are the same when redistributing OSPF into OSPF. Therefore, when a route is added into the routing table by AS X and a route already exists for the same destination by AS Y, it will not be added because the administrative distance is the same. This conflict can cause routes to take inter-AS paths when an intra-AS path exists. The solution in Cisco is to change the administrative distance on all external AS routes to a value higher than the default of 110.¹ Therefore, intra-AS paths will be preferred over inter-AS paths. It is important to note that this solution prevents all inter-AS paths when an intra-AS path exists, no matter what the cost is. This configuration may be considered unfavorable if the lowest cost path is always preferred or if the network topology contains many ASes.

Figure 7, shows that when redistribution is used the time between LSA installs outside the MANET is much higher than inside the MANET in a slower changing (less partitioning) network. This is because the routes to the prefixes in the MANET change, but there is still a route to the prefixes, so a new external LSA is not needed. As the network partitions more frequently, the count of LSA installs inside and outside the MANET approach equality because the routes to prefixes don't just change, but they are lost and gained. The benefit of this approach is that it only exposes changes in the MANET when partitions or merges in the MANET would affect the larger networks' ability to route to a MANET router. The disadvantage is that a suboptimal MBR may be chosen that causes a packet to take a high cost path to the router within the MANET.

Another strategy can be used in this scenario when the parallel radio links have dissimilar ranges. Emulation was performed when the MANET routers had high and low range wireless links. In this case, the LSAs from the MANET that impacted the larger network were minimal because the high range links enabled a stable route to all prefixes even when the low range links were changing rapidly. In fact, the LSAs in the larger network were merely from the LSA Refresh. Therefore, this technique can be used to protect the larger network from a highly changing MANET.

In the third scenario, summarization of the MANET prefixes is enabled by creating a tunnel between the MBRs over the backbone network. The tunnel uses links in the backbone but

¹http://www.cisco.com/en/US/tech/tk365/technologies_tech_note09186a00801069aa.shtml




Fig. 7. Larger network linked with a MANET using redistribution of Type 1 LSAs and static cost




Fig. 8. Larger network linked with a MANET using redistribution, summarization, and tunneling

appears in the control plane to be situated within the MANET. In this configuration, the network is able to sustain a MANET partition because if a packet is sent to the wrong MBR then it can be tunneled to a better MBR. Also, tunneling can be used to pick the MBR that gives the best path within the MANET if one is willing to sacrifice the per-packet overhead penalty of tunneling. Figure 8 shows that the time between installing an LSA is very large; much higher than in Figures 6 and 7. In fact, LSAs sent by the MBRs are not due to MANET changes at all. They are merely due to OSPF's design to refresh LSAs every LSRefreshTime seconds. Also, there are many fewer External LSAs in the larger network because the redistributed routes are summarized. The host routes are no longer being advertised individually. Table II provides a qualitative comparison of the three routing techniques to connect an OSPF-based MANET to a larger OSPF network.

V. SUMMARY

This paper has explored the use of several techniques for interconnecting a MANET running OSPF MANET extensions with a larger OSPF network. We first created a taxonomy of different approaches, and provided a rough bound on how many link state advertisements would be circulated in the OSPF flooding domain. We next described in detail the different approaches available, using routing techniques (OSPF hierarchy, cross-layer abstraction, BGP interconnection) and mobility management techniques (mobile mesh, readdressing). Using a mobile network emulator, we explored the performance benefit of partitioning the OSPF flooding domain via standard redistribution techniques, compared with the performance of flat OSPF/OSPF MANET routing (using a single

Configuration	Benefits	Drawbacks
Single Area	Lowest cost routes and good MANET entry points	MANET effects are completely exposed
AS Redistribution	Only MANET partitions and merges are exposed	Bad entry point to MANET and must redistribute all MANET routes (no summarization)
AS Redistribution and Summarization with Tunneling	Good MANET entry point and no MANET changes exposed	Must configure tunnels and suffer delay and bandwidth penalties from tunneling

TABLE II

COMPARISON OF DIFFERENT OSPF INTEGRATION TECHNIQUES.

flooding domain). We described why particular variants of OSPF redistribution are more favorable for allowing correct routing even in the face of partitions, while keeping link advertisements low when there is no partitioning. Finally, we explored the benefit of creating virtual links (tunnels) between MANET Border Routers and found the technique to perform the best in terms of minimizing the routing updates on the legacy network, while trading off some path optimality in the data plane. A more detailed exploration of tunneling architectures could be the subject of future research, as well as similar integration issues for multicast routing [14] and the extension of OSPF MANET with cross-layer techniques [15].

ACKNOWLEDGMENTS

This work was supported by Office of Naval Research (ONR) contract N00014-06-C-0023 and was performed in collaboration with the Naval Research Laboratory Information Technology Division. The authors would like to thank Joe Macker and Jeff Weston of NRL for technical collaboration, and Santanu Das (ONR program manager) and Jae H. Kim (Boeing PI and PM) for their support.

REFERENCES

- [1] S. Corson and J. Macker, "Mobile Ad hoc Networking (MANET): Routing Protocol Performance Issues and Evaluation Considerations," RFC 2501 (Informational), Jan. 1999. [Online]. Available: <http://www.ietf.org/rfc/rfc2501.txt>
- [2] J. Moy, "OSPF Version 2," RFC 2328 (Standard), Apr. 1998. [Online]. Available: <http://www.ietf.org/rfc/rfc2328.txt>
- [3] R. Coltun, D. Ferguson, and J. Moy, "OSPF for IPv6," RFC 2740 (Proposed Standard), Dec. 1999. [Online]. Available: <http://www.ietf.org/rfc/rfc2740.txt>
- [4] P. Spagnolo and T. Henderson, "Comparison of Proposed OSPF MANET Extensions," in *Proceedings - IEEE Military Communications Conference MILCOM*, vol. 2. IEEE, Oct. 2006.
- [5] R. Ogier and P. Spagnolo, "MANET Extension of OSPF using CDS Flooding," IETF, Internet-Draft (work in progress) draft-ogier-manet-ospf-extension-07, March 2006.
- [6] M. Chandra and A. Roy, "Extensions to OSPF to Support Mobile Ad Hoc Networking," IETF, Internet-Draft (work in progress) draft-chandra-ospf-manet-ext-04, January 2007.
- [7] C. A. Santivanez, R. Ramanathan, and I. Stavrakakis, "Making link-state routing scale for ad hoc networks," in *MobiHoc*. ACM, 2001, pp. 22–32.
- [8] J. Weinstein, J. Zavgren, B. Elliott, N. Rehn, and W. Passman, "Radio network routing apparatus," US Patent Number 6,977,937 B1, Dec. 2005. [Online]. Available: <http://patft.uspto.gov/netacgi/nph-Parser?patentnumber=6977937>
- [9] Y. Rekhter and T. Li, "A Border Gateway Protocol 4 (BGP-4)," RFC 1771 (Draft Standard), Mar. 1995, obsoleted by RFC 4271. [Online]. Available: <http://www.ietf.org/rfc/rfc1771.txt>

- [10] K. Grace, "Mobile Mesh Routing Protocol," IETF, Internet-Draft (work in progress) draft-grace-manet-mmrip-00, September 2000. [Online]. Available: http://www.mitre.org/work/tech_transfer/mobilemesh/draft-grace-manet-mmrip-00.txt
- [11] T. Henderson, "End-Host Mobility and Multihoming with the Host Identity Protocol," IETF, Internet-Draft (work in progress) draft-ietf-hip-mm-06, April 2007.
- [12] V. Devarapalli, R. Wakikawa, A. Petrescu, and P. Thubert, "Network Mobility (NEMO) Basic Support Protocol," RFC 3963 (Proposed Standard), Jan. 2005. [Online]. Available: <http://www.ietf.org/rfc/rfc3963.txt>
- [13] M. Zec and M. Mikuc, "Operating System Support for Integrated Network Emulation in IMUNES," *Proceedings of 1st Workshop on Operating System and Architectural Support for the on demand IT InfraStructure (ASPLOS-XI)*, Oct. 2004.
- [14] I. Chakeres, C. Danilov, T. Henderson, and J. Macker, "Connecting MANET Multicast," in *Proceedings – IEEE Military Communications Conference MILCOM*. IEEE, Oct. 2007.
- [15] G. Pei, P. Spagnolo, S. Bae, T. Henderson, and J. Kim, "Performance Improvements of OSPF MANET Extensions: A Cross Layer Approach," in *Proceedings – IEEE Military Communications Conference MILCOM*. IEEE, Oct. 2007.
- [16] *Proceedings of the Military Communications Conference (MILCOM), Orlando, FL, USA, October 29-31, 2007*. IEEE, 2007.

CONNECTING MANET MULTICAST

Ian D. Chakeres, Claudiu Danilov, Thomas R. Henderson, Joseph P. Macker[†]

Boeing Phantom Works
P.O. Box 3707, MC 7L-49
Seattle, WA 98124-2207

[†]Information Technology Division, Naval Research Laboratory
Washington, DC

Abstract—Connecting mobile ad-hoc networks (MANETs) to the legacy layer networks poses numerous challenges that have not been much explored. The internal topology of a MANET and external connectivity to a backbone network will likely vary greatly over time, and MANET nodes need to be able to communicate with nodes in other attached networks as changes occur. Keeping a MANET's multicast routing domain connected to other networks is not well understood and has been under-explored, especially when the MANET has multiple connections to the larger network. In this paper, we describe the challenges for attaching a MANET and delivering multicast between the MANET and other, higher tiered, legacy networks. Most importantly, we provide an efficient simple design solution that addresses the various multicast challenges. Additionally, we analyze the qualitative performance of different design elements and make recommendations for actual deployment.

I. INTRODUCTION

Multicast is the dominant form of traffic for today's tactical edge communications. Applications such as combat net radio and blue force tracking naturally have a one-to-many or many-to-many communications patterns, and it has been estimated that 70+% of tactical Internet-based traffic is multicast [1]. Although multicast dominates traffic, most mobile ad hoc network (MANET) [2] research has focused on unicast packet delivery. Furthermore, interconnection of MANETs with legacy IP networks is under explored and undefined, especially in regard to multicast.

While there are several multicast approaches used in legacy IP networks, this paper focuses on interconnecting MANET multicast with the Protocol Independent Multicast Sparse-Mode (PIM-SM) [3], a widely used multicast standard today. PIM-SM operates efficiently on the wired network media, but it is challenged by the MANET environment. Specifically, PIM-SM's tree structure, slow change and convergence time, and loop detection mechanisms make it not suitable for deployment over MANETs.

In MANET, Simplified Multicast Forwarding (SMF) [4] is an emerging standard for distributing multicast. SMF is simple, and robust to topology changes. SMF is efficient when used with a reduced relay set and can often outperform more complex multicast algorithms [5].

SMF and PIM-SM use very different approaches at the protocol level and interconnecting them is a challenge. At present, to the best of our knowledge, there are only static stopgap solutions to interconnecting the two routing regions.

In this paper, we explain the existing multicast structure of MANET networks and their attachment to legacy networks. We define the requirements for a capable multicast connection strategy, and specify the behavior of multicast MANET border routers (MBR) to fulfill these requirements. We follow up with a qualitative discussion of the various tradeoffs to support and manage multiple MBRs, and with a preliminary quantitative evaluation. Finally, we close with some conclusions about our approach.

II. BACKGROUND

A. IGMP and MLD

The Internet Group Membership Protocol (IGMP) [6] is the standard mechanism used by IPv4 hosts to communicate their group membership requirements to their attached routers. A derivative of IGMPv2 was defined as Multicast Listener Discovery (MLD) for IPv6 hosts [7]; in our discussions below, we describe operations in terms of IGMP, however, it equally applies to MLD. Routers generally solicit reports from hosts periodically using IGMP queries. Hosts may also issue IGMP leave messages when their membership changes. From the host perspective, after issuing IGMP messages their router(s) are responsible for delivering multicast to them.

B. SMF

SMF is a simple method of supporting and maintaining multicast forwarding within a MANET. When an IP multicast packet is received by a SMF node, if the packet is not a duplicate (and the packet passes the criteria for forwarding), the SMF node forwards the packet on its configured interfaces. The SMF node updates and maintains a duplicate packet detection (DPD) cache. The DPD cache ensures that packets are not forwarded multiple times, thereby avoiding forwarding loops. An example of sending a data packet through SMF is presented in Figure 1.

SMF nodes forwarding multicast form a loose mesh that is robust to topology changes. When using classical flooding or forwarding, there is no maintained state other than the DPD cache. SMF can be used in conjunction with a reduced relay set, reducing redundant or unneeded transmissions. For example, in the case of classic flooding presented in Figure 1, when the router with dotted transmission lines transmits, all its neighboring routers have already received the multicast packet. Using a reduced relay set has been shown to maintain the overall SMF robustness while improving its efficiency [8].

Ian D. Chakeres is now with Motorola Labs India.




Fig. 1. SMF Multicast. Packets are propagated from the originating host to all MANET routers.




Fig. 2. PIM RPF check. The dotted line indicates that the link is not used by the destination router to reach the source. RPF check fails on that link but succeeds on the valid reverse unicast path.




Fig. 3. A PIM-SM tree

SMF can be considered a dense-mode multicast protocol, since multicast is delivered to all the SMF nodes within a MANET. It does not include any group membership controls, unlike most other traditional dense-mode and sparse-mode protocols. SMF has been designed to allow additional filtering in its forwarding state, even dynamically. While it is possible to build group membership filters on top of the current SMF, the main purpose of the SMF protocol itself is to specify the forwarding mesh method without pruning.

C. PIM-SM

PIM-SM is the most widely deployed multicast routing protocol today. PIM-SM utilizes the unicast routing table to assist in multicast routing, performing *reverse path forward* (RPF) checks. The RPF check ensures that no multicast loops are formed. When a multicast packet is received, the PIM-SM node checks the unicast routing table to determine the interface used to reach the multicast packet's IP source address. If the route indicates that the source is reachable via the same interface that the multicast packet was received from, then the packet is said to have passed the RPF check. If the RPF check passes, then the packet is forwarded on the interfaces indicated by the multicast forwarding table. If the RPF check fails, the packet is not forwarded [3].

Since all PIM routers are normally within the same routing region, they will have a nearly synchronized view of each other's routing tables and should not encounter a RPF check failure during steady-state operation. Figure 2 shows a simple PIM network.

PIM-SM uses hello messaging to discover nearby PIM-SM routers. PIM-SM also has Join and Prune messages which manipulate the multicast forwarding tree. PIM-SM uses the

unicast routing table to help determine the interface to send Join and Prune messages. For a particular multicast group, the PIM-SM tree is rooted at a Rendezvous Point (RP). The RP for each multicast group is a well known IP address. Various mechanisms exist for statically or dynamically electing RPs [9]. For this paper, we assume RPs are statically configured and well-known.

When a host subscribes to a multicast group, its PIM-SM router issues a PIM Join message to the PIM-SM router nearer to the RP. Eventually a shortest-path tree is built back to the RP. At this point, if traffic is received at the RP, it will flow down the created PIM-SM tree. Figure 3 shows the multicast forwarding tree formed via PIM Join messages.

When a multicast source transmits a packet, the PIM-SM router on the host's link encapsulates the multicast packet in a PIM Register message. The PIM Register message is unicast to the RP, as shown in Figure 4. Once the RP receives the PIM Register message, it de-encapsulates the multicast packet and sends it down the PIM-SM tree, as shown in Figure 3. At this point multicast is being sent to the subscribed receivers.

If topology changes occur in the PIM-SM tree, PIM messaging will repair the tree. Repairing the tree may take on the order of tens of seconds, as PIM-SM is designed for static wired networks.

PIM-SM has various mechanisms to optimize performance, one of which is switching to native multicast. This can be done by the RP or a receiver's router at the edge of the PIM-SM tree.

Instead of encapsulating multicast traffic in PIM Register messages, the RP can build forwarding routes toward itself using PIM messaging. The advantage of native multicast messaging is avoiding encapsulation. The RP initiates this change and once the native forwarding routers are working,




Fig. 4. A PIM Register message

the RP issues a PIM Register-Stop message to stop a source's PIM-SM router from sending encapsulated packets. Thereafter, the RP must periodically issue PIM Register-Stop messages to keep the source's PIM-SM router from using encapsulation.

Switching to a tree rooted at the multicast source can improve efficiency within the PIM-SM domain. A subscribed PIM-SM router can initiate this change using PIM-SM messaging. Once the native multicast packets are being received via the route (interface) toward the multicast IP source, each PIM-SM router is free to prune links toward the RP if unneeded.

D. Related Work

a. Routers Participating in Multiple Multicast Protocols

During the development of PIM-SM, there was work on running multiple routing protocols and the mechanisms required to inter-operate [10]. This previous work has focused on the use case of a router fully participating in both PIM-SM and SMF (multiple multicast routing protocols), and is practical when MBRs are statically connected to a PIM-SM network. In contrast, we focus on the case where MBRs are chosen from the pool of MANET routers as those that happen to have external connectivity, regardless of where they are attached to the external network.

b. Homogeneity vs Heterogeneity

In this work, we discuss connecting network layer multicast. Network layer multicast forwarding allows the multicast to flow over multiple heterogeneous interface types.

There has been existing work on homogeneous networks where all links use the same technology, and in these networks multicast is generally handled below the network layer. The advantage of handling multicast below the network layer is that network layer multicast protocols may not need to be modified. The disadvantages are that it prevents heterogeneous interconnection, which is one of the major advantages of IP networks, makes additional layer-3 filtering harder, and results in higher overall overhead [2].

III. CONNECTING MANET MULTICAST

In our framework, we assume a network as shown in Figure 5. Hosts indicate their group membership to their attached router(s) using IGMP. Note that MANET routers may also act as hosts and subscribe themselves to multicast groups. SMF is used within the MANET to disseminate multicast packets. Within the backbone, PIM-SM is used as the multicast routing protocol. At the border of the SMF and PIM-SM regions

lie MANET Border Routers (MBRs). MBRs are responsible for handling the disparities between SMF and PIM-SM, and ultimately delivering multicast between both regions.

Our assumed design requirements are as follows. We must maintain existing host-router communication mechanisms, namely IGMP; hosts must remain unchanged and work inside a MANET. Second, the backbone network usually cannot be modified, since it will likely be composed of legacy routers (LRs) and protocols. We assume that the interconnection between the two multicast domains can be best accomplished at MBRs.

In addition, an acceptable solution must be able to handle multiple MBRs simultaneously. This requirement stems from the fact that a MANET must be capable of maintaining operation as its connectivity changes, within the MANET and to the backbone network. For example, if a MANET partitions into two MANETs and each partition still has a connected MBR, then we would like multicast to continue to flow properly into and out of the MANET.

We assume a link-state unicast routing protocol in the MANET. This assumption is not strictly required, but it is likely an operational need to have some unicast capability to supplement multicast. Our design will attempt to leverage this unicast routing protocol if it is available. We have the most familiarity with the MANET extensions for OSPF, and have developed an interface to allow SMF to leverage its CDS information.

Using OSPF provides a few additional pieces of information and mechanisms to coordinate multicast routing. Specifically, OSPF will provide all MRs information about the membership within the MANET; that is, each MR will learn information about all of the other MRs in the system. This information, leveraged by the multicast protocol, will allow us to significantly improve performance without increasing complexity. OSPF also allows additional information to be carried in its protocol messages [Opaque]. The additional information is carried in a type-length-value (TLV) format. We use TLVs to distribute information from MRs to MBRs.

Other routing protocols (e.g., OLSR) that provide membership information and the ability to add arbitrary TLVs could be used instead of OSPF without modification to our multicast strategy.




Fig. 5. The assumed network architecture

A. Ingress - Getting Multicast into a MANET

There are two main requirements to enable multicast traffic from selected groups to enter a MANET from a PIM-SM region. First, MBRs must indicate interest to the upstream PIM-SM routers. Second, they must forward received multicast datagrams from the PIM-facing interface to the MANET interface.

In our design, we consider that MBR act as an IGMP proxy for nodes within the MANET. This requires little configuration and state maintenance on the MBR, and it quickly engages the attached PIM-SM router. Figure 6 shows MANET ingress via multiple MBRs. Each MBR subscribes to the multicast on behalf of hosts within the MANET. To the attached PIM-SM routers, the MBRs appear to be hosts.

MBRs are also responsible for forwarding non-MANET-originated multicast traffic into the MANET. This design fulfills the multiple ingress MBR case without issue. If there are multiple ingress MBRs, SMF tagger ID will allow detection of the multiple gateways, and provide duplicate packet detection per gateway. If the MANET later partitions, each MANET connected to a MBR will continue to receive multicast traffic from the PIM-SM network.

There is one missing component to this design if group membership is not statically configured: how the MRs may notify the MBR of their interest to receive multicast traffic. To inform the MBRs of MANET multicast group membership, MRs might distribute their hosts' multicast group membership using a link-state advertisement such as an OSPF opaque LSA for group membership. This mechanism could be similar to MOSPF's [11] mechanism for distributing group membership. It would differ from MOSPF in that the SMF nodes do not use the information to form a multicast tree. In this paper we did not explore this path; instead, we use SMF forwarding to disseminate IGMP messages to the entire MANET. By distributing group membership, MBRs can learn MANETs' group membership and proxy IGMP for MANET nodes. The IGMP proxy operation will be similar to the one presented in [12], but differs slightly since it is based upon OSPF LSAs instead of IGMP messages. If a non-link-state unicast protocol were in use in the MANET, two possibilities would be to




Fig. 6. MANET Ingress




Fig. 7. MANET Egress

multicast disseminate the group membership information, or have the interested MRs tunnel their request directly to the MBRs.

If there are multiple ingress routers in a MANET and only one is required, this solution might be considered to place an unnecessary load on the PIM-SM region. In Section IV, we discuss reducing the number of ingress MBR and the tradeoffs with doing so. Section IV-A discusses further optimizations for multiple MBRs.

As a simple alternative, we experimented with forwarding IGMP messages through SMF to the entire MANET, thus guaranteeing that membership information reaches MBRs.

B. Egress - Getting Multicast out of a MANET

We assume that MBRs have a means to learn PIM rendezvous points (RPs) for PIM-SM operation. At the MBR, one can use PIM Register messages to egress multicast packets from the MANET and into the PIM-SM region. This design enables multiple egress MBR. It also handles MANET topology and connectivity changes seamlessly.

Figure 7 shows MANET egress via multiple MBRs. A source transmits a multicast and its MR transmits the packet into the SMF region. SMF forwards the multicast packet throughout the MANET. When a MBR receives the multicast, it encapsulates the packet in a PIM Register message and unicasts the message to the RP.

The drawback of this approach is that if there are multiple MBRs, it will result in duplicate multicast packets being delivered via multiple PIM Register messages. Subsequently, each of these duplicates will also travel down the PIM-SM tree. We discuss optimizations and tradeoffs to multiple MBR operation in Section IV.

IV. HANDLING MULTIPLE BORDER ROUTERS

If the entire MANET is connected through a single MBR, then that MBR represents the MANET to the external network, forwarding multicast traffic in and out of the MANET. However, in practice it is expected that multiple border routers will be deployed for improved resilience and avoiding congestion bottlenecks. In the simplest scenario, all MBRs will forward multicast traffic. This mechanism relies on the PIM-RP to drop

packets for which RPF check fails in case of egress traffic, and on SMF's duplicate packet detection to drop duplicate packets inside the MANET for ingress traffic. While this offers a working solution, it generates unnecessary duplication in the external traffic, leading to high overhead.

There are several possible improvements to the simple multicast connection strategy described above. The most pressing improvements revolve around coordination between multiple border routers to reduce traffic in the PIM-SM region. These improvements may come at the expense of redundancy.

A. Ingress - Reduced/Single Point of Entry

One path is sufficient to ingress multicast traffic into a MANET. That is, only one MBR needs to subscribe. Using only one MBR has the advantage of reducing traffic in the PIM-SM region. On the other hand, if we assume the MBRs communicate and coordinate to create a single ingress point, then we lose some redundancy in case the MANET partitions.

If a partition occurs, then a MBR must detect and subsequently subscribe to multicast traffic. The detection of partition and recovery time may be too large to advocate reducing the number of subscribing MBR. This cost can be large, as performing the PIM-SM Join process may take a relatively long period of time; on the order of tens of seconds.

To maximize the multicast ingress connectivity in the face of changing MANET topology and membership, all MBRs should proxy IGMP messages to their attached PIM-SM routers.

In networks with a large number of MBRs, the number of MBRs can be reduced by using a coordination mechanism. This mechanism could be as simple as choosing the N lowest ID MBR to participate, or using more complex election criteria/algorithm. For this solution to work, MBR need to identify themselves as MBR to other MBR; this information could be shared using the unicast routing protocol messages.

Using these operations if the MANET membership were to change (as indicated by the unicast routing protocol) other MBR in the same routing region would detect the changes and could subscribe or unsubscribe as configured.

Note that the proposed solution does not attempt to optimize traffic within the PIM-SM region. For example, PIM operation could be optimized by considering which MBR is close to the RP or particular multicast sources. To reduce traffic in the PIM-SM region and make intelligent MBR (multiple) ingress decisions, additional information about the PIM-SM topology would be required.

B. Egress - Reduced/Single Point of Exit

One path is sufficient to egress MANET-generated multicast traffic from a MANET. That is, only one MBR needs to issue PIM Register messages. Using only one MBR has the advantage of reducing traffic in the PIM-SM region by a large amount, as duplicate multicast will traverse down the entire PIM-SM tree for the group unnecessarily. On the other hand, if we assume MBR communicate and coordinate to create a single egress, then we lose some redundancy in case the MANET partitions.




Fig. 8. A first-hop MR issuing PIM Register messages directly to the RP, on behalf of an attached host

In our solution using OSPF-MANET as the unicast routing protocol (or other routing protocol that provides MANET membership information), we recommend that MBR (or MR) coordinate to choose a single node that encapsulates multicast in PIM Register messages. This coordination could be as simple as using the MBR with the lowest ID or could take additional factors into consideration.¹ When a MANET partitions, the unicast routing protocol will disseminate routing information about the change, enabling MBRs to determine whether they should become active and issue PIM Register messages on behalf of nodes within their partition.

This solution would eliminate redundant PIM Register messages, and more importantly, it would stop redundant multicast packets from being passed down the PIM-SM tree.

C. Egress - First-hop MANET Router Encapsulation

If there are no multicast subscribers for a particular group in the MANET and there is a multicast source for this group in the MANET, the first-hop MR can issue PIM Register messages on its behalf. This optimization eliminates unneeded flooding of this multicast data within the SMF region. This optimization would require that MR monitor group membership information within the MANET.

Figure 8 shows a MANET router encapsulating its source's multicast in a PIM Register message and delivering it directly to the RP without MANET wide dissemination.

In this case there is only a single egress, and there would be no multiple egress challenge.

D. Egress - Single Point of Exit using Native Multicast

To avoid encapsulation and the use of PIM Register messages in the PIM-SM region, a single MBR could utilize native IP multicast forwarding. The difficulty in attaining this single point of exit is that the MBRs are likely not informed well enough to know whether their forwarding of multicast will pass the PIM-SM RPF check. That is, due to PIM-SM's RPF check - some of the possible egress border routers may not be able to inject native multicast traffic.

¹To inform MBR in a MANET about each other, an additional TLV or flag could be added to OSPF advertisements to indicate that a particular MR is a MBR.




Fig. 9. MANET PIM RPF difficulty

Figure 2 provides an example of PIM's RPF check. In this section we provide Figure 9 to describe a MANET's difficulty. Since the MBRs are participating in multiple routing regions simultaneously, they may or may not know the routing information for the PIM routing region. If the wrong egress route is chosen, and the PIM RPF check fails, then those packets will be lost. If a single egress is used, this would result in non-delivery of multicast packets. Also, there is no indication when packets are dropped due to RPF check failure.

An alternative would be to have all MBRs participate in native multicast, and forward multicast packets to their closest PIM router. The RPF check will ensure that packets are not duplicated on the path from the RP to the end-hosts, although PIM routers may still forward duplicate packets towards the RP. In our preliminary implementation, presented in Section V, we chose this approach. The drawback of using native multicast at the MBRs is that node movement within the MANET that determines partitioning between MBRs can take on order of seconds to tens of seconds to update the PIM-SM tree forwarding, as we will see in Section V.

E. Ingress - DPD Marking

When sources (or first-hop routers) don't set the IP ID or IP option for DPD properly, this function should be performed by ingress MBR. Unfortunately, several problems can arise when DPD sequencing is done by multiple ingress MBRs.

To counteract problems with multiple ingress MBR DPD sequencing, we propose including the Tagger (MBR) ID inside an added IP option for DPD. This modification would stop multiple ingress border routers' packets from colliding in the DPD ID space. The disadvantage of this approach is that SMF will flood packets with identical payload if they have different DPD keys.

This problem could be counteracted by several means. MBRs could passively detect multiple points of ingress (from reception of other MBRs' tagged packets). Another solution would be a MBR election, like lowest ID; this mechanism would suffice to stop multiple ingress of the same multicast packet and different DPD marking. This method would also solve the multiple ingress gateway coordination issue, but it would impose additional state maintenance on the MBR, and perhaps delay delivery if a partition occurs.

V. PROTOCOL EXPERIMENTATION

In order to test the feasibility of our approach we emulated a topology consisting of nine MANET nodes connected to a PIM network through two MBRs. We used a network emulator developed at Boeing, based on the open-source




Fig. 10. Experiment topology




Fig. 11. Ingress to Mobile: Latency.

IMUNES [13] project, which provides an entire network stack virtualization and topology control inside a single FreeBSD machine. The emulated topology is presented in Figure 10. The nine MANET nodes were running Quagga [14] OSPFv3 with MANET extensions for unicast routing, and SMF for multicast. The PIM routers were running the XORP [15] routing software with PIM-SM, and were using OSPFv2 as the unicast routing protocol. IGMP messages issued by nodes inside the MANET were disseminated via SMF throughout the entire MANET, and the border routers were acting as multicast proxies, forwarding multicast traffic in and out of the MANET and IGMP messages out of the MANET to the PIM routers.

A mobile node inside the MANET was moving following the trajectory shown in Figure 10. All MANET links were set to 20ms delay, and the wired link between PIM-RP and PIM2 was set to 100ms delay. The use of these latency values in the emulation helps to graphically illustrate the number of hops taken by multicast packets in Figures 11 - 13.

In a first experiment we sent multicast traffic from Host 1 to the Mobile node, at a rate of 5 packets of 1024 bytes per second. The packet latency results are shown in Figure 11. Lost packets are represented as having latency zero. We can see that, as the Mobile node was moving, the recorded packet latency changed with the number of hops traversed inside the MANET. When the Mobile node was partitioned together with the MBR 2 side of the network (around sequence number 450




Fig. 12. Egress to Host 1: Latency.

in Figure 11), packets sent by Host 1 had to traverse the high latency link between PIM-RP and PIM 2, adding about 100ms in latency. During the experiment, only one packet was lost, at sequence number 632.

In a second experiment we sent multicast traffic from the Mobile node, and had the two hosts attached to the PIM network join as receivers. The recorded packet latency is shown in Figure 12 and Figure 13, for Host 1 and Host 2, respectively. We can see that, in general, hosts attached to the PIM network received packets continuously, regardless of the Mobile node movement inside the MANET. The only interruption in connectivity appeared for Host 1 at around sequence number 400, when the MANET partitioned by moving the Mobile node from the network connected to MBR 1 to the network connected to MBR 2, causing 14 packets to be lost in a row (about 3 seconds). Throughout the experiment, Host 1 experienced one additional loss at sequence number 487, and Host 2 experienced a total of 3 packet losses.

Our preliminary implementation had both MBRs egress multicast traffic. As a result, duplicate packets were sent by the two MBRs within the PIM domain. However, inside the MANET, SMF's duplicate packet detection was eliminating duplicates introduced by the two MBRs (although possibly not reducing duplicates of the initial transmissions from co-located MBRs). Packet duplication inside the PIM network can cause high overhead, especially when the number of MBRs is large, and therefore solutions that reduce the number of egress MBRs (see Section IV) should be deployed.

VI. CONCLUSIONS

In this paper, we propose mechanisms to connect MANET multicast to legacy networks. Specifically, we define how MANET Border Routers need to behave to transport SMF MANET multicast traffic to a PIM-SM region and vice versa. The design also allows multiple MBRs with little or no coordination among them. Our solution is also robust to MANET partition and changing MBR connectivity. The design allows for various optimizations.

We recommend that a MANET take advantage of multiple ingress MBRs to handle network partitions quickly. We also recommend that only one node performs PIM Register messaging on behalf of its MANET, and we describe how this can easily be accomplished when using OSPF.




Fig. 13. Egress to Host 2: Latency.

ACKNOWLEDGMENTS

This work was supported by Office of Naval Research (ONR) contract N00014-06-C-0023 and was performed in collaboration with the Naval Research Laboratory Information Technology Division. The authors would like to thank Santanu Das (ONR program manager) and Jae H. Kim (Boeing PI and PM) for their support.

REFERENCES

- [1] "Control-Based Mobile Ad-Hoc Networking Program Industry Day Presentation," http://www.darpa.mil/sto/solicitations/CBMANET/proposers_day.htm.
- [2] I. D. Chakeres and C. E. Perkins, "Dynamic MANET On-demand (DYMO) Routing Protocol," *IETF Internet Draft, draft-ietf-manet-dymo-04.txt*, March 2005, (Work in Progress).
- [3] D. Estrin, D. Farinacci, A. Helmy, D. Thaler, S. Deering, M. Handley, V. Jacobson, C. Liu, P. Sharma, and L. Wei, "Protocol Independent Multicast-Sparse Mode (PIM-SM): Protocol Specification," *RFC 2362*, June 1998.
- [4] J. Macker, J. Dean, and W. Chao, "Simplified multicast forwarding in mobile ad hoc networks," in *Military Communications Conference, 2004. MILCOM 2004. IEEE*, October 2004.
- [5] K. Obraczka, G. Tsodik, and K. Viswanath, "Exploring Mesh- and Tree-Based Multicast Routing Protocols for MANETs," in *IEEE Transactions on Mobile Computing*, Vol. 5, No. 1, pp. 28-42, January 2006.
- [6] W. Fenner, "Internet group management protocol, version 2," IETF, RFC 2236, Nov. 1997. [Online]. Available: <ftp://ftp.isi.edu/in-notes/rfc2236.txt>
- [7] S. Deering, W. Fenner, and B. Haberman, "Multicast Listener Discovery (MLD) for IPv6," *RFC 2710*, October 1999.
- [8] T. Kunz, "Multicast versus broadcast in a manet," in *ADHOC-NOW*, 2004, pp. 14-27.
- [9] J. Mukherjee, R. Atwood, "Rendezvous point relocation in protocol independent multicast - sparse mode," in *International Conference on Telecommunications, 2003*, March 2003.
- [10] D. Thaler, "Interoperability Rules for Multicast Routing Protocols," *RFC 2715*, October 1999.
- [11] J. Moy, "Multicast extensions to OSPF," IETF, RFC 1584, Mar. 1994. [Online]. Available: <ftp://ftp.isi.edu/in-notes/rfc1584.txt>
- [12] B. Fenner, H. He, B. Haberman, and H. Sandick, "Internet Group Management Protocol (IGMP) Multicast Listener Discovery (MLD)-Based Multicast Forwarding," *RFC 4605*, August 2006.
- [13] M. Zec, "Implementing a clonable network stack in the freebsd kernel," in *Proceedings of the 2003 USENIX Annual Technical Conference*, June 2003.
- [14] "Quagga Routing Software Suite," <http://www.quagga.net>.
- [15] "XORP: the eXtensible Open Router Platform," <http://www.xorp.org>.

Boeing RANGE Software

Installation Guide and User's Manual

Authors:

Philip A. Spagnolo

Claudiu Danilov

Thomas Goff

Ian D. Chakeres

Thomas R. Henderson

Copyright © 2006-08, The Boeing Company.

This software was developed under funding supplied by ONR contract N00014-06-C-0023.

Table of Contents

1	Introduction	1
1.1	Software directory overview	1
1.2	Licensing	2
1.2.1	OSPF-MANET	2
1.2.2	NRL SMF	2
1.2.3	XORP	2
2	Prerequisites	4
2.1	Patching the source code	4
2.1.1	OSPF-MANET	4
2.1.2	SMF	4
2.1.3	XORP	4
3	OSPFv3 Address Families	5
3.1	Overview	5
3.2	Configuring OSPF-AF	5
3.3	Running OSPF-AF	5
3.4	Open Issues	6
4	OSPF-MANET	7
4.1	Overview	7
4.1.1	Draft compliance	7
4.1.2	Contributors	7
4.2	Protocol Operation	7
4.3	Building OSPF-MANET	7
4.4	Configuring OSPF-MANET	8
4.5	Running OSPF-MANET	9
4.6	Use with Address Families	9
4.6.1	Redistribution between OSPFv2 and OSPFv3 MANET	9
4.7	OSPF-MANET Configuration Examples	10
5	SMF	11
6	XORP	12
6.1	PointToMultipoint	12
6.2	PIM-DM Interface	12

1 Introduction

This manual describes installation and usage of software developed under Boeing's Robust Airborne NetworkinG Extensions (RANGE) program.

There are four sets of software patches.

- Patches to enable OSPF-MANET functionality in quagga OSPFv3
- Patch to NRL SMF to enable it to work with XORP
- Patch to XORP OSPFv2 to enable the PointToMultiPoint interface type
- Patch that provides an initial implementation of a PIM-Dense Mode (PIM-DM) interface for XORP

1.1 Software directory overview

The top-level directory contains three sub-directories:

- quagga/
- xorp/
- smf/

and a PDF (range.pdf) of this document.

The quagga subdirectory contains the following:

- quagga-0.99.9.tar.gz (unmodified release)
- quagga.pdf (unmodified manual for version 0.99.9)
- Boeing's OSPFv3 Extensions patch: quagga-0.99.9.ospfv3-extensions.patch
- Boeing's OSPFv3 Address Families patch: quagga-0.99.9.ospfv3-addressfamilies.patch
- Boeing's OSPFv3 MANET Designated Routers (MDR) patch: quagga-0.99.9.ospfv3-manetmdr.patch

The xorp subdirectory contains the following:

- xorp-1.4.tar.gz (unmodified release)
- xorp_user_manual.pdf (unmodified manual for version 1.4)
- Boeing's OSPF PointToMultiPoint patch: xorp-1.4-ptmpospf-011508.patch
- Boeing's PIM-DM Interface: xorp-1.4-dm.patch

The smf subdirectory contains the following:

- src_nrlsmf-1.1b1.tar (unmodified release)
- Boeing's SMF patch to enable XORP integration: smf-1.1b1.patch

1.2 Licensing

All software has been cleared for public release as open source, but it is important for users to understand the licensing associated with each piece of software.

1.2.1 OSPF-MANET

Boeing's OSPF-MANET related software is a derivative work of the Quagga routing suite, which is licensed under the GNU General Public License (GPL) version 2. Therefore, the quagga-0.99.9 patches are provided under GPL version 2 (Copyright 2008 Boeing):

```
This program is free software; you can redistribute it and/or
modify it under the terms of the GNU General Public License
as published by the Free Software Foundation; either version 2
of the License, or (at your option) any later version.
```

```
This program is distributed in the hope that it will be useful,
but WITHOUT ANY WARRANTY; without even the implied warranty of
MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. See the
GNU General Public License for more details.
```

```
You should have received a copy of the GNU General Public License
along with this program; if not, write to the Free Software
Foundation, Inc., 51 Franklin Street, Fifth Floor, Boston, MA
02110-1301, USA.
```

1.2.2 NRL SMF

NRL SMF extensions are provided by NRL with no licensing terms. Boeing's modifications to NRL SMF are provided under the same terms (Copyright 2008 Boeing).

1.2.3 XORP

XORP routing software is provided under the following license

Copyright (c) 2001-2008 International Computer Science Institute

```
Permission is hereby granted, free of charge, to any person obtaining a
copy of this software and associated documentation files (the "Software"),
to deal in the Software without restriction, including without limitation
the rights to use, copy, modify, merge, publish, distribute, sublicense,
and/or sell copies of the Software, and to permit persons to whom the
Software is furnished to do so, subject to the following conditions:
```

```
The above copyright notice and this permission notice shall be included in
all copies or substantial portions of the Software.
```

```
The names and trademarks of copyright holders may not be used in
advertising or publicity pertaining to the software without specific
prior permission. Title to copyright in this software and any associated
```

documentation will at all times remain with the copyright holders.

THE SOFTWARE IS PROVIDED "AS IS", WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO THE WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT. IN NO EVENT SHALL THE AUTHORS OR COPYRIGHT HOLDERS BE LIABLE FOR ANY CLAIM, DAMAGES OR OTHER LIABILITY, WHETHER IN AN ACTION OF CONTRACT, TORT OR OTHERWISE, ARISING FROM, OUT OF OR IN CONNECTION WITH THE SOFTWARE OR THE USE OR OTHER DEALINGS IN THE SOFTWARE.

Boeing's modifications and extensions to XORP software are provided under the same licensing terms (Copyright 2008 Boeing).

2 Prerequisites

The system requires a relatively modern Linux distribution (with GNU development tools such as the gcc compiler and make) and related packages. We have tested this with Fedora Core 5 and 6 distributions. Root privileges are needed.

As for host hardware requirements, either a Linux computer with an Ethernet connection, or a Linux virtual machine on some other (e.g., Windows) operating system, should work.

2.1 Patching the source code

The software is provided as a patch because you may want to apply the patch to some other base release of the software (and it may still work).

2.1.1 OSPF-MANET

There are three patches provided. Each patch provides incrementally more functionality. **You should select and apply only one of the three patches!**

First, unpack quagga:

```
tar xvfz quagga-0.99.9.tar.gz
```

Next, select a patch to apply:

1. Boeing's OSPFv3 Extensions patch: quagga-0.99.9.ospfv3-extensions.patch
2. Boeing's OSPFv3 Address Families patch: quagga-0.99.9.ospfv3-addressfamilies.patch
3. Boeing's OSPFv3 MANET Designated Routers (MDR) patch: quagga-0.99.9.ospfv3-manetmdr.patch

Patch 1 provides just some OSPFv3 extensions, intended for general OSPF enhancement. This is probably not interesting to RANGE users. Patch 2 (Address Families) provides a patch corresponding to the Address Families extension for OSPFv3, but no MANET software. Patch 3 provides the code in patches 1 and 2, and adds the OSPF MANET MDR code. If in doubt, apply patch 3:

```
patch -p0 < quagga-0.99.9.ospfv3-manetmdr.patch
```

2.1.2 SMF

The below lines will patch the unmodified SMF distribution:

```
tar xvf src-nrlsmf-1.1b1.tar
patch -p0 < smf-1.1b1.patch
```

2.1.3 XORP

The below lines will patch the unmodified XORP distribution for PointToMultipoint:

```
tar xvfz xorp-1.4.tar.gz
patch -p0 < xorpt-1.4-ptmpospf-011508.patch
```

The below lines will patch the unmodified XORP distribution for the PIM-DM interface:

```
tar xvfz xorp-1.4.tar.gz
patch -p0 < xorp-1.4-dm.patch
```

The PIM-DM patch already includes the patch for PointToMultipoint.

3 OSPFv3 Address Families

The following text describes the implementation of a mechanism for supporting multiple address families in OSPFv3 using multiple instances. It maps an address family (AF) to an OSPFv3 instance using the Instance ID field in the OSPFv3 packet header. This approach is fairly simple and minimizes extensions to OSPFv3 for supporting multiple AFs.

This implementation also enables OSPF MANET to support IPv4 routing (next chapter).

3.1 Overview

For now, please see Section 2 of [draft-ietf-ospf-af-alt-05](#). Support of Address Families (AF) in OSPFv3 is supported according to version 5 of the above draft.

3.2 Configuring OSPF-AF

Address Families can be configured in one of two ways.

1. Add the following lines to the ospf6d.conf file

```
interface <ifname>
ipv6 ospf6 instance-id <0-255>
```

2. From the vtysh or telnet terminal type:

```
> conf t
> interface <ifname>
> ipv6 ospf6 instance-id <0-255>
> exit
> exit
```

The value of the instance ID should be in one of the four ranges below. The most common ranges are 0 to 31 for unicast IPv6 routing (standard OSPFv3) and 64 to 95 for IPv4 unicast routing.

Instance ID # 0	- # 31	IPv6 unicast AF
Instance ID # 32	- # 63	IPv6 multicast AF
Instance ID # 64	- # 95	IPv4 unicast AF
Instance ID # 96	- # 127	IPv4 multicast AF
Instance ID # 128	- # 255	Unassigned

NOTE: The instance-id must be the same on all interfaces. Different Address Families cannot be used within the same ospf6d process. The router will fail if different ranges are used.

3.3 Running OSPF-AF

From a vtysh or telnet terminal type:

```
> show ipv6 ospf6 route
```

This should display the OSPFv3 routes. If IPv4 AFs are used then the route will appear as an IPv6 route with zeros before the IPv4 route. Next, type the following command for IPv4 or IPv6

```
> show ip route  
> show ipv6 route
```

The entries with the "*" are going to be installed in the kernel routing table. If these tables are correct then the kernel routing table should be correct.

3.4 Open Issues

Enable different AFs to run in the same ospf6d instance. This would require IETF draft changes and a separation of LSAs within the database.

Known Issue: if instance IDs are not consistent on the interfaces then routing will fail.

4 OSPF-MANET

OSPF-MANET is a modification of OSPF version 3 (IPv6) for use in mobile ad hoc networks (MANETs). OSPF for IPv6 is described in RFC2740.

This chapter is a supplement to the main quagga (<http://www.quagga.net>) documentation. It describes the implementation, functionality, and usage of OSPF-MANET and related extensions.

4.1 Overview

OSPF-MANET can be built for typical quagga usage as a standalone router, for support in virtual machines such as IMUNES, and within a discrete-event network simulator.

OSPF-MANET is defined for IPv6 (OSPFv3). With the addition of what is known as the *Address Families* patch, an instance of OSPF-MANET can also be run to build IPv4 routes. Note that to get both IPv4 and IPv6 routing, two instances of OSPFv3 must be running, as presently defined by the draft standard.

OSPF-MANET is distributed as a series of patches against a mainline quagga distribution. The Boeing server is located at

<http://hipserver.mct.phantomworks.org/ietf/ospf/>.

4.1.1 Draft compliance

[draft-ogier-manet-ospf-extension-07](#)

4.1.2 Contributors

This OSPF-MANET software is the product of a number of individuals, including:

- Jeff Ahrenholz
- Claudiu Danilov
- Tom Henderson
- Jeff Meegan
- Richard Ogier
- Gary Pei
- Phil Spagnolo

and collaboration with Naval Research Laboratory.

4.2 Protocol Operation

For now, please see

<http://hipserver.mct.phantomworks.org/ietf/ospf/milcom06.pdf> or Section 2 of [draft-ogier-manet-ospf-extension-10](#)

4.3 Building OSPF-MANET

To build quagga as standalone router run:

```
./configure --enable-user=root --enable-group=root --enable-vtysh \
--with-cflags=-ggdb
make
make install
```

4.4 Configuring OSPF-MANET

OSPF-MANET can be configured in one of two ways: command line interface (CLI) or config file (ospf6d.conf). In either case, you must install a zebra.conf and ospf6d.conf file in /usr/local/etc/.

- CLI: run configuration commands in vtysh or telnet
- put configuration commands in zebra.conf and ospf6d.conf

Here are the configuration commands added during the development of OSPF-MANET MDR.

```
* router ospf6
** router minls-interval <0-65535> : Minimum time between LSA
origination.
** router minls-arrival <0-65535> : Minimum time between LSA
reception.

* interface <ifname> : Select the interface to configure
** ipv6 ospf6 network (broadcast|non-broadcast|point-to-multipoint|
point-to-point|loopback|manet-designated-router)
*** broadcast: Specify OSPF6 broadcast multi-access network
*** non-broadcast: Specify OSPF6 NBMA network
*** point-to-multipoint: Specify OSPF6 point-to-multipoint network
*** point-to-point: Specify OSPF6 point-to-point network
*** loopback: Specify OSPF6 loopback
*** manet-designated-router: Specify OSPF6 manet-designated-router (MDR)
network
** ipv6 ospf6 flood-delay <1-65535> : Time in msec to coalesce LSAs before
sending
** ipv6 ospf6 jitter <1-65535> : Time in msec to jitter sending of
all ospf6 packets
** ipv6 ospf6 ackinterval <1-65535> : Interval of time to coalesce acks
** ipv6 ospf6 backupwaitinterval <1-65535> : Interval of time for MBDRs to
wait before flooding
** ipv6 ospf6 diffhellos : Enable differential hellos
** ipv6 ospf6 twohoprefresh <1-65535> : When using differential Hellos,
full Hellos are sent every TwoHopRefresh Hellos.
** ipv6 ospf6 hellorepeatcount <1-65535> : Total hellos in succession that
cannot be missed using diff hellos
** ipv6 ospf6 adjacencyconnectivity (unconnected|biconnected|fully) :
Level of adjacencies between neighbors
*** unconnected: The set of adjacencies forms a (uni)connected graph.
*** biconnected: The set of adjacencies forms a biconnected graph.
```

```

*** fullyconnected: Adjacency reduction is not used, the router becomes
adjacent with all of its neighbors.
** ipv6 ospf6 lsafullness (minlsa|mincostlsa|mincost2lsa|mdrfulllsa|fulllsa):
Choose the OSPFv3 interface type
*** minlsa: Specify min size LSAs (only adjacent neighbors)
*** mincostlsa: Specify partial LSAs for min-hop routing
*** mincost2lsa: Specify partial LSAs for two min-hop routing paths
*** mdrfulllsa: Specify full LSAs from MDR/MBDRs
*** fulllsa: Specify full LSAs (all routable neighbors)

```

4.5 Running OSPF-MANET

Run the following commands for the command prompt:

```

/usr/local/sbin/zebra -d
/usr/local/sbin/ospf6d -d

```

To verify OSPF-MANET is running, from a vtysh or telnet terminal type:

```
> show ipv6 ospf6 route
```

This should display the OSPFv3 routes. If IPv4 AFs are used then the route will appear as an IPv6 route with zeros before the IPv4 route. Next, type the following command for IPv4 or IPv6

```

> show ip route
> show ipv6 route

```

The entries with the "*" are going to be installed in the kernel routingtable. If these tables are correct then the kernel routing table should be correct.

4.6 Use with Address Families

To use OSPF MANET to carry IPv4 prefix information, one may enable it with the following configuration.

In the interface description section, define an instance-id such that $64 < \text{instance-id} < 96$. For instance:

```

interface ath0
...
ipv6 ospf6 instance-id 65
...

```

Then, in the router definition section, describe networks to be associated with OSPF MANET.

```

router ospf6
router-id 10.1.0.1
interface ath0 area 0.0.0.0

```

4.6.1 Redistribution between OSPFv2 and OSPFv3 MANET

(to be completed)

4.7 OSPF-MANET Configuration Examples

Here is an example of an interface declaration of an OSPF-MANET interface, from the `ospf6d.conf` file.

```
interface ath0
    ipv6 ospf6 priority 1
    ipv6 ospf6 transmit-delay 1
    ipv6 ospf6 instance-id 65
    ipv6 ospf6 ifmtu 1500
    ipv6 ospf6 cost 1
    ipv6 ospf6 hello-interval 2
    ipv6 ospf6 dead-interval 6
    ipv6 ospf6 retransmit-interval 5
    ipv6 ospf6 network manet-designated-router
    ipv6 ospf6 ackinterval 1800
    ipv6 ospf6 diffhellos
    ipv6 ospf6 backupwaitinterval 2000
    ipv6 ospf6 twohoprefresh 3
    ipv6 ospf6 hellorepeatcount 3
    ipv6 ospf6 adjacencyconnectivity biconnected
    ipv6 ospf6 lsafullness mdrfulllsa
    ipv6 ospf6 flood-delay 100
!
```

The below router declaration example tells quagga to run OSPF-MANET on interface `ath0` and to redistribute OSPF and connected networks.

```
router ospf6
    router-id 10.1.0.1
    interface ath0 area 0.0.0.0
    redistribute ospf
    redistribute connected
!
```

5 SMF

The enclosed patch provides a few small changes to SMF to allow it to be run with XORP. One change is that XORP adds an interface to the machine, but then SMF complains that there are too many. Another change has to do with the duplicate packet detection logic, when XORP and SMF coexist on the machine.

Another change forces IGMP messages to be flooded throughout the SMF region. This is contrary to the design of IGMP which is meant to only be sent over a single hop. This change was necessary, so that PIM gateways could see MANET routers (that may be multiple hops away) that have subscribed to a multicast group.

In addition to the change to the SMF code, the TTL on IGMP messages must be increased. This can be done through the use of firewall rules, such as this example.

```
iptables -t mangle -I OUTPUT 1 -p IGMP -j TTL --ttl-set 7
```

Note that this command should be run after starting XORP and SMF in order to insert the iptables rule in front of other rules that may be set by SMF.

6 XORP

This section provides information about using the Boeing patches for XORP.

6.1 PointToMultipoint

This is a small patch to enable the OSPFv2 Point-To-Multipoint mode. Point-To-Multipoint is described in RFC 2328; it basically allows one to aggregate multiple point-to-point OSPF relationships on a single interface. It is the architectural basis for the OSPF MANET interfaces.

The interface type can be used like other XORP OSPF interface types Point-to-Multipoint (PTMP) OSPF can be used in one of two ways: unicast or multicast. The unicast version requires the user to configure the neighboring routers manually during configuration. The multicast version dynamically finds and adds neighbors. The PTMP patch changes XORP's PTMP implementation to better fit a wireless broadcast environment.

6.2 PIM-DM Interface

The XORP code changes to allow PIM-DM functionality are based on the existing PIM-SM code, in the pimsm4 directory. The PIM-DM interface is not a PIM-DM implementation of RFC 3973. It is an interface that supports PIM-DM functionality between a PIM-DM network and an SMF network.

Most code changes deal with modifying the state machine for data forwarding from "default off" in PIM-SM, to a "default on" approach in PIM-DM, and with triggering appropriate control message exchanges, as required by the PIM-DM specification (Assert, Prune and Graft). The current implementation handles responds to all PIM-DM messages issued by neighboring routers, and issues Assert messages as needed. However, in the current implementation, the router does not issue Prune and Graft messages based on the dynamic group membership available behind interfaces configured for IGMP, even though it does block the corresponding multicast traffic when there are no subscribers for a certain group. As a consequence, when used as a MANET border gateway, our router forwards traffic towards the MANET correctly, based on the MANET membership, but does not reduce multicast traffic in the external network when packets are not needed in the MANET. Multicast traffic originated inside the MANET is correctly blocked at the gateways when it is not needed in the external network, or forwarded by only one of the gateways when external subscribers exist.

Configuring and using the PIM-DM router is the same as for a PIM-SM router, with the only difference that the Rendezvous Point has to be set as the address of a local interface on the router, instead of a single interface of a unique, potentially remote router selected as the RP.